

Simulated evolution of protein-protein interaction networks with realistic topology

Jack Peterson,^{1,*} Steve Pressé,² Kristin S. Peterson,³ and Ken A. Dill⁴

¹*Biophysics Graduate Group, University of California, San Francisco, CA 94158*

²*Department of Pharmaceutical Chemistry, University of California, San Francisco, CA 94158*

³*Department of Forest Ecosystems and Society, Oregon State University, Corvallis, OR 97331*

⁴*Laufer Center for Physical and Quantitative Biology, Stony Brook University, NY 11794*

We model the evolution of eukaryotic protein-protein interaction (PPI) networks. In our model, PPI networks evolve by two known biological mechanisms: (1) Gene duplication, which is followed by rapid diversification of duplicate interactions. (2) Neofunctionalization, in which a mutation leads to a new interaction with some other protein. Since many interactions are due to simple surface compatibility, we hypothesize there is an increased likelihood of interacting with other proteins in the target protein's neighborhood. We find good agreement of the model on 10 different network properties compared to high-confidence experimental PPI networks in yeast, fruit flies, and humans. Key findings are: (1) PPI networks evolve modular structures, with no need to invoke particular selection pressures. (2) Proteins in cells have on average about 6 degrees of separation, similar to some social networks, such as human-communication and actor networks. (3) Unlike social networks, which have a shrinking diameter (degree of maximum separation) over time, PPI networks are predicted to grow in diameter. (4) The model indicates that evolutionarily old proteins should have higher connectivities and be more centrally embedded in their networks. This suggests a way in which present-day proteomics data could provide insights into biological evolution.

We are interested in the evolution of protein-protein interaction (PPI) networks. PPI network evolution accompanies cellular evolution, and may be important for processes such as the emergence of antibiotic resistance in bacteria [1, 2], the growth of cancer cells [3], and biological speciation [4–6]. In recent years, increasingly large volumes of experimental PPI data have become available [7–10], and a variety of computational techniques have been created to process and analyze these data [11–18]. Although these techniques are diverse, and the experimental data are noisy [19], a general picture emerging from these studies is that the evolutionary pressures shaping protein networks are deeply interlinked with the networks' topology [20]. Our aim here is to construct a minimal model of PPI network evolution which accurately captures a broad panel of topological properties.

In this work, we describe an evolutionary model for eukaryotic PPI networks. In our model, protein networks evolve by two known biological mechanisms: (1) a gene can duplicate, putting one copy under new selective pressures that allow it to establish new relationships to other proteins in the cell, and (2) a protein undergoes a mutation that causes it to develop new binding or new functional relationships with existing proteins. In addition, we allow for the possibility that once a mutated protein develops a new relationship with another protein (called the target), the mutant protein can also more readily establish relationships with other proteins in the target's neighborhood. One goal is to see if random changes based on these mechanisms could generate networks with the properties of present-day PPI networks. Another goal is then to draw inferences about the evolutionary histories of PPI networks.

THE MODEL

We represent a PPI network as a graph. Each node on the graph represents one protein. A link (edge) between two nodes represents a physical interaction between the two corresponding proteins. The links are undirected and unweighted. To model the evolution of the PPI graph, we simulate a series of steps in time. At time t , one protein in the network is subjected to either a gene duplication or a neofunctionalizing mutation, leading to an altered network by time $t + \Delta t$. We refer to this model as the DUNE (DUPLICATION & NEOfunctionalization) model.

Gene duplication

One mechanism by which PPI networks change is gene duplication (DU) [21–23]. In DU, an existing gene is copied, creating a new, identical gene. In our model, duplications occur at a rate d , which is assumed to be constant for each organism. All genes are accessible to duplication, with equal likelihood. For simplicity, we assume that one gene codes for one protein. One of the copies continues to perform the same biological function and remains under the same selective pressures as before. The other copy is superfluous, since it is no longer essential for the functioning of the cell [24].

The superfluous copy of a protein/gene is under less selective pressure; it is free to lose its previous function and to develop some other function within the cell. Due to this reduced selective pressure, further mutations to the superfluous protein are more readily accepted, including those that would otherwise have been harmful to the organism [25, 26]. Hence, a superfluous protein diverges rapidly after its DU event [27, 28]. This well-known process is referred to as the *post-duplication di-*

* jack@tinybike.net




FIG. 1. **DUNE model flowchart.** At each time step, the simulated network undergoes a duplication or neofunctionalization event. Red nodes/links indicate nodes/links that have been created by duplication during the current time step. Green links indicate links that have been created by neofunctionalization during the current time step. A dashed line indicates a duplicated link that has been deleted during the post-duplication divergence. Only 3 neighbors are shown for the assimilation mechanism; however, the actual simulations included up to 20th neighbors. The simulated network evolves until its number of links (K) meets or exceeds the number of links in the data (K_{data}).

vergence. Following [29], we assume that the link of each such superfluous protein/gene to its former neighbors is deleted with probability ϕ . The post-duplication divergence tends to be fast; for simplicity, we assume the divergence occurs within the same time step as the DU. The divergence is asymmetric [30, 31]: one of the proteins diversifies rapidly, while the other protein retains its prior activity. We delete links from the original or the duplicate with equal probability because the proteins are identical. As discussed in the supporting information (SI), this is closely related to the idea of *subfunctionalization*, where divergence freely occurs until redundancy is eliminated. In our model, ϕ is an adjustable parameter.

In many cases, the post-duplication divergence results in a protein which has lost all its links. These ‘orphan’ proteins correspond to silenced or deleted genes in our model. As discussed below, our model predicts that the gene loss rate should be slightly higher than the duplication rate in yeast, and slightly lower in flies and humans.

We simulate a gene duplication event at time t as follows:

- 1a) Duplicate a randomly-chosen gene with probability $d\Delta t$.
- 2a) Choose either the original (50%) or duplicate (50%), and delete each of its links with probability ϕ .
- 3a) Move on to the next time interval, time $t + \Delta t$.

Neofunctionalization

Our model also takes into account that DNA can be changed by random mutations. Most such mutations do not lead to changes in the PPI network structure. However, some protein mutations lead to new interactions with some other protein (which we call the *target protein*). The formation of a novel interaction is called a *neofunctionalization* (NE) event. NE refers to the creation of new interactions, not to the disappearance of old ones. Functional deletions tend to be deleterious to organisms [32]. We do not account for loss-of-function mutations (link deletions) except during post-duplication divergence because damaged alleles will, in general, be eliminated by purifying selection. In our model, NE mutations occur at a rate μ , which is assumed to be constant. All proteins are equally likely to be mutated.

How does the mutated protein choose a target protein to which it links? We define a probability q that any protein in the network is selected for receiving the new link from the mutant protein. To account for the possibility of homodimerization, the mutated protein may also link to itself [24, 33]. Random choice dictates that $q = 1/N$ (see SI).

Many PPI’s are driven by a simple geometric compatibility between the surfaces of the proteins [34]. The simplest example is the case of PPI’s between flat, hydrophobic surfaces [35], a type of interaction which is very common [36]. These PPI’s have a simple planar interface, and the binding sites on the individual proteins are geometrically quite similar to one another. One consequence of these similar-surface interactions is that if protein A can bind to proteins B and C, then there is a greater-than-random chance that B and C will interact with each other. We refer to this property as *transitivity*: if A binds B, and A binds C, then B binds C. The number of triangles in the PPI network should correlate roughly with transitivity. As discussed below, the number of triangles (as quantified by the global clustering coefficient) is about 45 times higher in real PPI networks than in an equally-dense random graph. This suggests that transitivity is quite common in PPI networks. Another source of transitivity is gene duplication. If A binds B, then A is copied to create a duplicate protein A’, then A’ will (initially) also bind B. If A interacts with A’, then a triangle exists. However, duplication is unlikely to be the primary source of transitivity; recent evidence shows that, due to the post-duplication divergence, duplicates tend to participate in fewer triangles than other proteins [37].

A concrete example of transitivity is provided by the

evolution of the retinoic acid receptor (RAR), an example of neofunctionalization which has been characterized in detail [38]. Three paralogs of RAR exist in vertebrates (RAR α , β , and γ), as a result of an ancient duplication. The interaction profiles of these proteins are quite different. Previous work indicates that RAR β retained the role of the ancestral RAR [38], while RAR α and γ evolved new functionality. RAR α has several interactions not found in RAR β . RAR α has novel interactions with a histone deacetylase (HDAC3) as well as seven of HDAC3's nearest-neighbors (HDAC4, MBD1, Q15959, NRIP1, Q59FP9, NR2E3, GATA2). None of these interactions are found in RAR β . The probability that all of these novel interactions were created independently is very low. RAR α has 65 known PPI's and HDAC3 has 83, and the present-day size of the human PPI network is a little over 3000 proteins. Therefore, the chance of RAR α randomly evolving novel interactions with 7 of HDAC3's neighbors is less than 1 in a billion. This strongly suggests that when a protein evolves an interaction to a target, it has a greater-than-random chance of also linking to other, neighboring proteins.

How do similar-surface interactions affect the evolution of PPI networks? First, consider how an interaction triangle would form. Suppose proteins A and B bind due to physically similar binding sites. Protein X mutates and evolves the capacity to bind A. There is a reasonable chance that X has a surface which is similar to both A and B. If so, protein X is likely to also bind to B, forming a triangle. Denote the probability that two proteins interact due to a simple binding site similarity by a . The probability that A binds B (and X binds A) in this manner is a . Assuming these probabilities are identical and independent, the probability that X binds B is a^2 .

So far, we have discussed transitivity as it affects the PPI's in which protein A is directly involved (A's first-neighbors). We now introduce a third protein to the above example, resulting in a chain of interactions: protein A binds B, B binds C, but C does not bind A. Protein X mutates and gains an interaction with A (with probability a^2). What is the probability that X will also bind C? The probability that B binds C due to surface similarity is a . Thus, X will bind C (A's second-neighbor) with probability a^3 . In general, the probability that X will bind one of A's j^{th} neighbors is a^{j+1} . We refer to this process as *assimilation*, and the 'assimilation parameter' a is a constant which varies between species. As discussed in SI, it is primarily multiple-partner proteins which bind to their partners at different times and/or locations which are affected by this process; consequently, at most one link is created by assimilation at the first-neighbor level, second-neighbor level, etc. Assimilation is assumed to act on a much shorter time scale than duplication and neofunctionalization; in our model, it is instantaneous.

Our hypothesized assimilation mechanism makes several predictions that could be tested experimentally: (1) the probability of a protein assimilating into a new pathway should be a^2 (at the first-neighbor level), a^3 (at the

second-neighbor level), and so on, where a is a constant which varies between species; (2) weak, nonspecific binding and planar interfaces should be overrepresented in interaction triangles (and longer cycles) between non-duplicate proteins; (3) competitive inhibitors should be overrepresented in interaction triangles; and (4) domain shuffling should be associated with assimilation. (See SI for discussion of (3) and (4).)

We simulate a neofunctionalization event at time t as follows:

- 1b) Mutate a randomly-chosen gene with probability $\mu\Delta t$.
- 2b) Link to a randomly-chosen target protein.
- 3b) Add a second link to one of the target's first-neighbor proteins, chosen randomly, with probability a^2 .
- 4b) Add a link to one of the target's second-neighbor proteins, with probability a^3 , etc.
- 5b) Move on to the next time interval, time $t + \Delta t$.

Model simulation and parameters

A flowchart of how PPI networks evolve in our model is shown in Figure 1. To simulate the network's evolution, one of the two mechanisms above is used at each time step, using [39]. We call each possible time series a *trajectory*. We begin each trajectory starting from two proteins sharing a link (the simplest configuration that is still technically a network). Each simulated trajectory ends when the model network has grown to have the same total number of links, K , as found in the experimental data, K_{data} . Here, we perform sets of simulations for three different organisms: yeast (*Saccharomyces cerevisiae*), fruit flies (*Drosophila melanogaster*), and humans (*Homo sapiens*). Because evolution is stochastic, there are different possible trajectories, even for identical starting conditions and parameters. We simulated 50 trajectories for each organism. Our figures show the me-

TABLE I. Network sizes and model parameters

	N_{data}	K_{data}	d	μ	ϕ	a
Yeast	2170	3819	0.01	7.86×10^{-4}	0.555	0.690
Fly	878	1140	0.0014	5.89×10^{-4}	0.866	0.546
Human	3165	5547	0.0037	7.62×10^{-4}	0.652	0.727

N and K are the numbers of proteins and links, respectively. (K_{data} is used to stop the simulation. N_{data} is not used as a constraint.) d and μ have units of per gene per million years (Myr). ϕ and a are probabilities (unitless). K_{data} and d are constraints from the data, while μ , ϕ , and a are adjustable parameters. We used Monte Carlo simulations to optimize the parameter values, by minimizing the total symmetric mean absolute percentage error values of the simulated versus the experimental data (see SI). Our values of μ are substantially lower than d because μ is the rate of mutations leading to the creation of a new PPI (rather than being a simple mutation rate, which would be much higher).




FIG. 2. **Degree centrality.** Degree (k) distributions in human (green), yeast (blue), and fly (red). Heavy lines are the median values from 50 simulations, and light lines are results of individual simulations. Points represent high-confidence empirical data for each organism (see Methods). Unless otherwise noted, color coding is the same in all plots. Quantitative comparisons between simulation and experiment (for DUNE and several other models) are detailed in SI.

dian values of each feature as a heavy line, and individual trajectories as light lines.

For a given data set, the number of links (K_{data}) is known. We estimate the duplication rate d from literature values. There have been several empirical estimates of duplication rates, mostly falling within an order of magnitude of each other [27, 40–45]. We averaged together the literature values to estimate d for each species (Table I).

The quantity μ is not as well known. Its value relative to d has been the topic of considerable debate [24, 46–48]. Although, in principle, μ is a measurable quantity, it has proven difficult to obtain an accurate value, in part because the fixation rate of neofunctionalized alleles varies with population size [49, 50]. In the absence of a consensus order-of-magnitude estimate, in our model, we treat μ as a fitting parameter. Consistent with the findings of [51] and [46], our best-fit values of μ are within an order of magnitude of each other for yeast, fruit fly, and human networks. Best-fit parameter values are given in Table I.

RESULTS

Present-day network topology

One test of an evolutionary model is its predictions for present-day PPI network topologies. Current large-scale PPI data sets have a high level of noise, resulting in significant problems with false positives and negatives

[19, 54]. To mitigate this, we compare only to ‘high-confidence’ experimental PPI network data gathered in small-scale experiments (see Methods). We computed 10 topological features, quantifying various static and dynamic aspects of the networks’ global and local structures: degree, closeness, eigenvalues, betweenness, modularity, diameter, error tolerance, largest component size, clustering coefficients, and assortativity. 8 of these properties are described below (see SI for others).

The *degree* k of a node is the number of links connected to it. For protein networks, a protein’s degree is the number of proteins with which it has direct interactions. Some proteins interact with few other proteins, while other proteins (called ‘hubs’) interact with many other proteins. Previous work indicates that hubs have structural and functional characteristics that distinguish them from non-hubs, such as increased proportion of disordered surface residues and repetitive domain structures [56]. The high degree of a protein hub could indicate that protein has unusual biological significance [57]. The network’s overall link density is described by its mean degree, $\langle k \rangle$ (Table II). The *degree distribution* $p(k)$ is the probability that a protein will have k links. PPI networks have a few hub proteins and many relatively isolated proteins. The heavy tail of the degree distribution shows that PPI networks have significantly more hubs than random networks have. Simulated and experimental degree distributions are compared in Figure 2. (For quantitative comparisons, see SI.)

Component refers to a set of reachable proteins. If any protein is reachable from any other protein (by hopping from neighbor to neighbor), then the network only has one component. If there is no path leading from protein A to B, then A and B are in different components. The

	Q	D	f_1	$\langle C \rangle$	$\langle k \rangle$
Yeast data	0.75	15	0.89	0.09	3.65
DUNE	0.74(7)	17(6)	0.8(1)	0.041(9)	4.0(8)
Vázquez	0.80(4)	21(5)	0.2(1)	0.045(5)	2.6(4)
Berg	0.518(4)	12.0(7)	0.990(3)	0.0027(9)	4.10(3)
RG	0.910(3)	36(3)	0.987(6)	0.475(8)	5.31(8)
MpK	0.58(6)	24(5)	1.000(2)	0.08(3)	4.4(6)
ER	0.588(8)	13.0(9)	0.995(2)	0.002(1)	3.5(6)
Fly data	0.86	23	0.73	0.10	2.93
DUNE	0.82(2)	20(2)	0.81(3)	0.09(1)	2.36(9)
Human data	0.75	15	0.88	0.08	3.69
DUNE	0.74(6)	17(2)	0.88(4)	0.09(1)	3.7(4)

TABLE II. **Comparison of network features.** Modularity Q , diameter D , fraction of nodes in the largest component f_1 , global clustering coefficient $\langle C \rangle$, and $\langle k \rangle$ is the average degree of proteins the largest component. ‘Data’ is the empirical data, ‘DUNE’ is the model described here, ‘Vázquez’ is the duplication-only model of [29], ‘Berg’ is the link dynamics model [52], ‘RG’ is random geometric [53], ‘MpK’ is the physical desolvation model presented in [54], and ‘ER’ is an Erdős-Rényi random graph [55]. Simulated values are the median (\pm standard deviation) over 50 simulations. (See SI for details of each model’s setup and optimization.)




FIG. 3. **Closeness centrality.** (A) Closeness (ℓ) distributions in human (green), yeast (blue), and fly (red). Heavy lines are the median values from 50 simulations, and light lines are results of individual simulations. (B) Examples of networks with low average closeness $\langle \ell \rangle = 0.06$ (top; each node is generally far away from most other nodes because there are no ‘short cuts’) and high average closeness $\langle \ell \rangle = 0.28$ (bottom; the random connections allow each node to be only a short distance from the other nodes). Note that both networks pictured here have the same number of nodes ($N = 100$) and roughly the same average degree (top: $\langle k \rangle = 4$, bottom: $\langle k \rangle = 3.7$).

fraction of nodes in the largest component (f_1) is a measure of network fragmentation (Table II and Figure S2). Note that, although silent genes (proteins with no links) exist in real systems, these genes do not appear in data sets consisting only of PPIs. Therefore, calculations of f_1 for all models exclude orphan proteins (proteins with $k = 0$).

Gene loss, the silencing or deletion of genes, is known to play an important role in evolution. The loss of a func-

tioning gene will damage an organism, making the gene loss unlikely to be passed on. The exception is if the gene is redundant. Consistent with this reasoning, evidence suggests that many gene loss events are losses of one copy of a duplicated gene [30, 58]. Although empirical estimates of the gene loss rate varied considerably, a consistent finding across several studies is that the rates of gene duplication and loss are of the same order-of-magnitude [27, 41, 44]. This broad picture is in good agreement with our model. In our model, a gene is con-

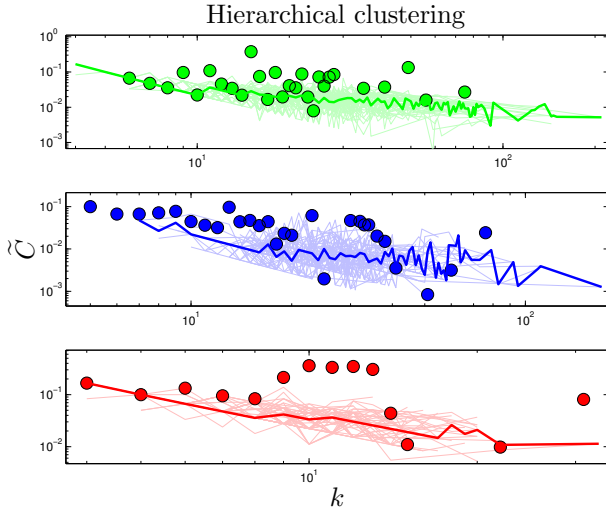


FIG. 4. **Hierarchical clustering.** Median clustering coefficient vs. degree in human (green), yeast (blue), and fly (red). Heavy lines are the median values from 50 simulations, and light lines are results of individual simulations.

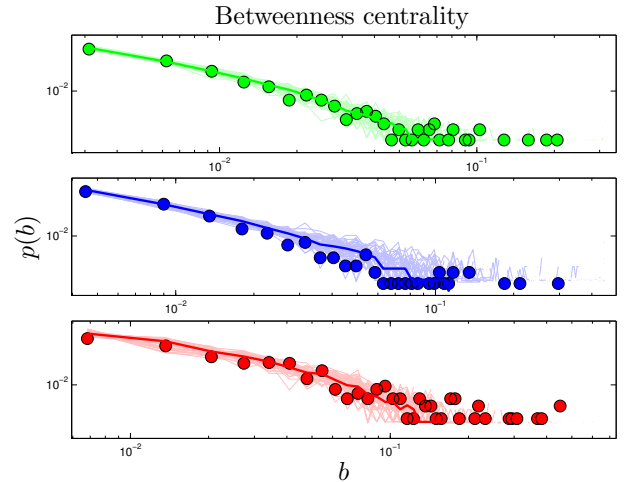


FIG. 5. **Betweenness centrality.** Betweenness (b) distributions in human (green), yeast (blue), and fly (red). Heavy lines are the median values from 50 simulations, and light lines are results of individual simulations.

sidered lost when it has degree zero. Our model predicts that the ratio of orphan to non-orphan proteins is 1.6 ± 0.4 in yeast, 0.58 ± 0.06 in flies, and 0.67 ± 0.09 in humans. The gene loss rate has been previously estimated to be about half the duplication rate in both flies and humans [27, 44], consistent with our model's prediction.

The *distance* between nodes i and j is defined as the number of node-to-node steps that it takes along the shortest path to get from node i to j . The *closeness centrality* of a node i , ℓ_i , is the inverse of the average distance from node i to all other nodes in the same component. The *diameter*, D , of a network is the longest distance in the network. Simulated closeness distributions are compared to experiments in Figure 3. Interestingly, proteins have about ‘six degrees of separation’, similar to social networks [59, 60]. The closeness distributions $p(\ell)$ have peaks around $1/\ell \approx 5 - 7$.

Another property of a network is its *modularity* [61]. Networks are modular if they have high densities of links (defining regions called modules), connected by lower densities of links (between modules). One way to quantify the extent of modular organization in a network is to compute the modularity index, Q [62, 63]:

$$Q \equiv \frac{1}{K} \sum_{i,j} \left(A_{ij} - \frac{k_i k_j}{K} \right) \delta(u_i, u_j), \quad (1)$$

where k_i and k_j are the degrees of nodes i and j , u_i and u_j denote the modules to which nodes i and j belong, $\delta(u_i, u_j) = 1$ if $u_i = u_j$ and $\delta(u_i, u_j) = 0$ otherwise, and $A_{ij} = 1$ if nodes i and j share a link, and $A_{ij} = 0$ otherwise. Q quantifies the difference between the actual within-module link density to the expected link density in a randomly connected network. Q ranges between -1




FIG. 6. **Betweenness vs. degree.** Shown are median betweenness vs. degree values in human (green), yeast (blue), and fly (red). Heavy lines are the median values from 50 simulations, and light lines are results of individual simulations.




FIG. 7. **Assortativity.** Median nearest-neighbor degree vs. degree in human (green), yeast (blue), and fly (red). Heavy lines are the median values from 50 simulations, and light lines are results of individual simulations.

and 1; positive values of Q indicate that the number of links within modules is greater than random. The numerical value of Q required for a network to be considered ‘modular’ depends on the number of nodes and links and method of computation. To calibrate baseline Q values given our particular network data, we used the null model described in [64]. Our non-modular baseline values are $Q = 0.603$ for the human PPI net, $Q = 0.590$ for yeast, and $Q = 0.722$ for flies (see SI). As shown in Table II, PPI networks are highly modular, and our simulated Q values are in good agreement with those of experimental data.

The *clustering coefficient*, C_i , for a protein i , is a measure of mutual connectivity of the neighbors of protein i . C_i is defined as the ratio of the actual number of links between neighbors of protein i to the maximum possible number of links between them,

$$C_i = \frac{\# \text{ edges between neighbors of node } i}{k_i(k_i - 1)}. \quad (2)$$

In a PPI network, clustering is thought to reflect the high likelihood that proteins of similar function are mutually connected [65]. The average (or global) clustering coefficient, $\langle C \rangle$, quantifies the extent of clustering in the network as a whole. As shown in Table II, PPI networks have large global clustering coefficient values; the yeast PPI network, for example, has a value of $\langle C \rangle$ which is 45 times higher than that of a random graph of equivalent link density. In flies and humans, our simulated networks have $\langle C \rangle$ values in excellent agreement with the data; in yeast, our predicted value is slightly low.

A network is said to be ‘hierarchically clustered’ if the clustering coefficient and degree obey a power-law relation, $C \sim k^{-\xi}$ [66] (Figure 4), indicating that nodes are

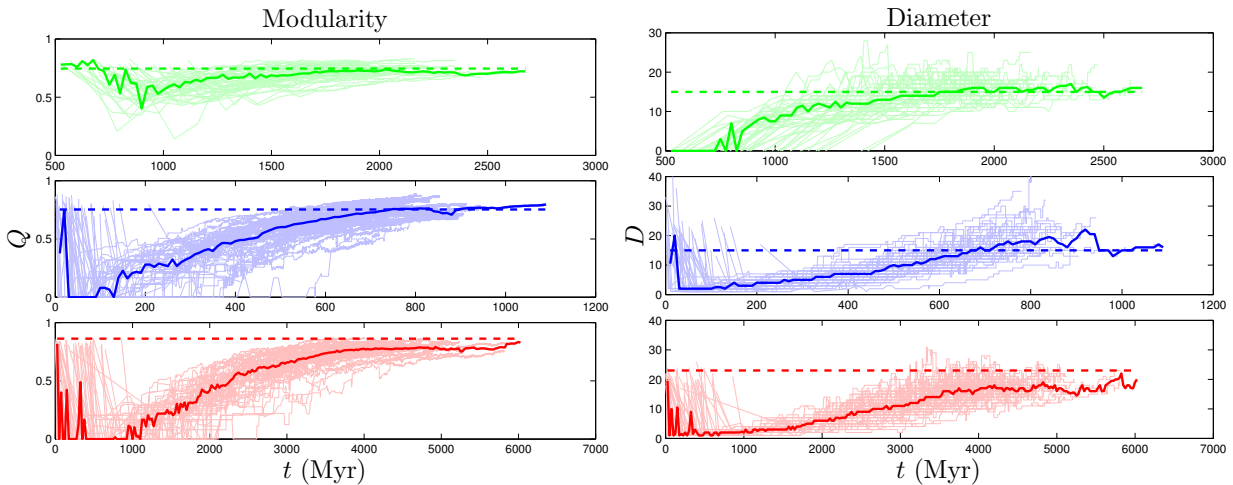


FIG. 8. **Modularity and diameter.** (A) Modularity Q and (B) diameter D are both predicted to grow with time in human (green), yeast (blue), and fly (red). Light lines indicate the evolutionary trajectories of 50 individual simulations, and the heavy line is the median value. The modularity and diameter of the empirical data are shown as dashed horizontal lines. Time traces occasionally do not start at $t = 0$ because these simulations spend the first few time steps in a completely disconnected state, so the dynamical quantities are undefined. (See Figure S2 for other dynamical plots.)

organized into small-scale modules, and the small-scale modules are in turn organized into larger-scale modules following the same pattern [67]. By plotting each node's clustering coefficient against its degree, we observed a trend consistent with hierarchical clustering, although data in the tail is very limited.

The *betweenness* of a node measures the extent to which it 'bridges' between different modules. *Betweenness centrality*, b_i , is defined as:

$$b_i \equiv \frac{\# \text{ shortest paths passing through node } i}{\# \text{ total shortest paths}}. \quad (3)$$

Betweenness has been proposed as a uniquely functionally-relevant metric for PPI networks because it relates local and global topology. It has been argued that knocking out a protein that has high betweenness may be more lethal to an organism than knocking out a protein of high degree [68]. Betweenness distributions are shown in Figure 5.

If a network's well-connected nodes are mostly attached to poorly-connected nodes, the network is called *disassortative*. A simple way to quantify disassortativity is by determining the median degree of a protein's neighbors (n) as a function of its degree (k). Previous work has found that yeast networks are disassortative [64]. It has been argued that disassortativity is an essential feature of PPI network evolution, and recent modeling efforts have heavily emphasized this feature [69, 70]. However, it was noted by [71] that disassortativity may simply be an artifact of the yeast two-hybrid technique, and [72] pointed out that this trend is quite different among different yeast datasets, and in some cases is completely reversed, resulting in *assortative* mixing, where high degree proteins prefer to link to other high-degree proteins.

As shown in Figure 7 and Table S1, the empirical data shows no evidence of disassortativity in flies or humans, and even the trend in yeast is quite weak. This conclusion is based solely on analysis of the empirical data, and casts further doubt on the role of disassortative mixing in PPI network evolution.

Some higher-order features of the network are simply a result of its degree sequence, and other features might be important in their own right. As discussed in [64], it is possible to isolate the effects of the degree sequence by 'rewiring' (detaching then reattaching links) the network at random, subject to the restriction that the degree sequence must be preserved. If a property contains extra information about the network's structure, then it should be different in the rewired network. On the other hand, if the network is rewired many times, and the property is always the same, then it is likely to just be a result of the degree sequence. We randomly rewired the empirical network $4K_{\text{data}}$ times.¹ As expected, modularity is decreased by random rewiring. Upon rewiring, we find $Q = 0.603 \pm 0.002$ in humans, $Q = 0.590 \pm 0.003$ in yeast, and $Q = 0.722 \pm 0.007$ in flies (median \pm standard deviation from 50 repeats of the rewiring algorithm). Rewiring also shrinks the diameters of PPI networks to $D = 13 \pm 0.9$ in humans, $D = 12 \pm 1.0$ in yeast, and $D = 15 \pm 1.1$ in flies. These results suggest that these features contain important structural information about the network, and are not merely consequences of the degree sequence.

¹ Network rewiring was carried out using a script downloaded from <http://www.cmth.bnl.gov/~maslov/matlab.htm>.

One reason we are interested in calculating Q and D is simply to check that the values are comparable between the simulated and experimental networks. However, on a more qualitative level, we would also like to have some idea of what the threshold is for a network to be considered ‘modular’ or ‘small-diameter’. The rewired Q and D values are useful because these features are dependent on the size of the network (number of nodes N and links K); given an identical network construction method, Q and D will generally be different in sparse versus dense networks. We use these Q and D values as baseline values with which the experimental and simulated networks can be compared; we considered Q and D values differing from the rewired values by more than a standard deviation to be significantly different.

Comparisons of simulated and experimental eigenvalue spectra and error tolerance curves are shown in SI (Figures S1 and S4). As discussed in SI, the various per-node network properties we have analyzed are largely uncorrelated (Figure S6).

Evolutionary trajectories

We now consider the question of how PPI networks evolve in time. The present-day networks show a rich-get-richer structure: PPI networks tend to have both more well-connected nodes and more poorly connected nodes than random networks have. In our model, the rich-get-richer property has two bases: duplication and assimilation. The equal duplication chance per protein means the probability for a protein with k links to acquire a new link via duplication of one of its interaction partners is proportional to k . Likewise, the probability of a protein to receive a link from the first-neighbor assimilation probability a is proportional to its degree k . ‘Rich’ proteins get richer because the probability of acquiring new links rises with the number of existing links.

First, we discuss two dynamical quantities for which experimental evidence exists: the rate of gene loss, and the relation between a protein’s age and its centrality. Gene losses in our model correspond to ‘orphan’ proteins which have no interactions with other proteins. As shown in Figure S2, the fraction of orphan proteins grows quickly at first, then levels off. This is consistent with the findings of [44]: in humans, while the overall duplication rate is higher than the loss rate, when only data from the past 200 Myr are considered, the loss rate is slightly higher than the duplication rate. In our model, after the initial rapid expansion, the rate of gene loss stabilizes relative to the duplication rate.

We define the ‘age’ of a protein in our simulation according to the order in which proteins were added to the network. Our model shows that a protein’s age correlates with certain network properties. Consistent with earlier work [73–76], we find that older proteins tend to be more highly connected. We plotted the ‘age index’ of a protein (the time step at which the protein was introduced) ver-




FIG. 9. **Gene ontology.** Shown are GO-slim profiles for the 100 oldest and 100 youngest proteins, as measured by S -value, in the yeast PPI network.




FIG. 10. **Older proteins are more central.** Simulations of a protein’s age index (time since introduction into the network) vs. degree (k), betweenness (b), and closeness (ℓ) centrality, for human (green), yeast (blue), and fly (red). The oldest proteins are on the *left* in this figure, and the proteins get younger moving to the right. There is an approximately monotonic increase in centrality with age.

sus its centrality scores. As shown in Figure 10, the age index negatively correlates with degree, betweenness, and closeness centralities: older proteins tend to be more central than younger proteins. Figure 10 shows our model’s prediction that a protein’s age correlates with degree, betweenness, and closeness centrality. We confirmed this prediction by following the evolutionary trajectories of individual proteins (Figure S3). These results are consistent with the eigenvalue-based aging method described in [76] (Figure S5). Phylogenetic protein age estimates indicate that older proteins tend to have a higher degree [73, 76], which our model correctly predicts. Interestingly, the eigenvalue-based scores are only modestly correlated with other centrality scores (0.36 degree, 0.47 betweenness, and 0.10 closeness correlations). Using the eigenvalue method in tandem with our centrality-based method could provide stronger age-discriminating power for PPI networks than either method alone.

The correlation between centrality and age suggests that static properties of present-day networks may be used to estimate relative protein ages. Suppose each normalized centrality score ($k' \equiv k / \max(k)$, $\ell' \equiv \ell / \max(\ell)$, $b' \equiv b / \max(b)$) represents a coordinate in a 3-D ‘centrality space’. We can then define a composite centrality score (S) as $S^2 \equiv (k')^2 + (\ell')^2 + (b')^2$.

Do older proteins typically have different functions than newer proteins? We classified *S. cerevisiae* proteins using the GO-slim gene ontology system in the Saccharomyces Genome Database. As shown in Figure 9, GO-slim enrichment profiles were somewhat different between the oldest and youngest proteins (as measured by their S values). Several categories which were more enriched for the oldest proteins were the cell cycle, stress response, cytoskeletal and cell membrane organization, whereas younger proteins were overrepresented in several metabolic processes. Overall, the differences were not dramatic, suggesting that cellular processes generally require both central and non-central proteins to function.

Consistent with this, ancient proteins tend to be centrally located with modules, as their betweenness values gradually decline over time (Figure S3). The roughly linear relation between degree and betweenness also suggests that ancient proteins do not occupy structurally ‘special’ positions within the network, such as stitching together separate modules (Table S1 and Figure 6). This may indicate that modules tend to accumulate around the most ancient proteins, which act as a sort of nucleus. Thus, ancient proteins are involved in all kinds of pathways, because they have each nucleated their own pathway.

In contrast to the two dynamical quantities discussed so far, most structural properties of PPI networks have only been measured for the present-day network. Although our model accurately reproduces the present-day values of these quantities, there is no direct evidence that the simulated trajectories are correct; rather, these are predictions of our model. Figure 8 shows that both modularity Q and diameter D increase with time. These are not predictions that can be tested yet for biological systems, since there is no time-resolved data yet available for PPI evolution. Time-resolved data is only currently available for various social networks (links to websites, co-authorship networks, etc.). Interestingly, the diameters of social networks are found to shrink over time [77]. Our model predicts that PPI networks differ from these social networks in that their diameters grow over time. In addition to Q and D , we tracked the evolutionary trajectories of several other quantities: the evolution of the global clustering coefficient, the rate of signal propagation, the size of the largest connected component (Figure S2), as well as betweenness and degree values for individual nodes (Figure S3). See SI for details.

The number of ‘extra’ links created by each assimilation event should independent of N . Proteins with multiple interaction partners may interact with their partners simultaneously (so-called ‘party hubs’) or at different times/locations (‘date hubs’) [78]. Since they bind

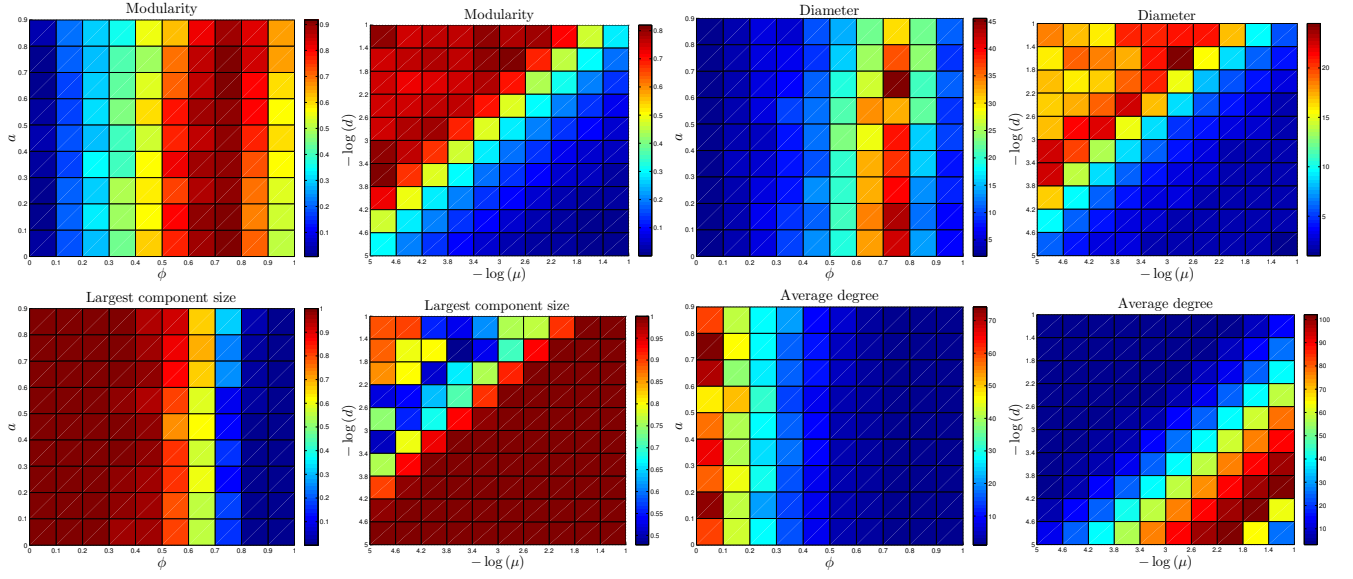


FIG. 11. **Sensitivity analysis.** Heat maps represent median values for 10 simulations per parameter combination of the yeast network. Left: ϕ and a are varied, d and μ values are kept fixed. Right: d and μ varied, ϕ and a kept fixed.

to several partners simultaneously, party hubs typically have multiple binding sites, each specific to one binding partner [79]. This specificity suggests that a protein which evolved the capability to bind to a party hub would be unlikely to undergo assimilation. By contrast, the binding sites of date hubs are often disordered regions which are able to form transient interactions with multiple partners [80, 81]. If a protein evolves the capability to bind to a date hub, it is likely to share the physical characteristics of the hub's neighbors, leading to assimilation. However, to avoid competition for the same binding site, the interaction partners of date hubs tend not to be coexpressed [79]. One consequence of this is that assimilating proteins will likely only bind one of the target protein's neighbors – whichever neighbor happens to be present at that time and place. Although the capability to bind to the hub protein's other neighbors may initially be present, these will presumably remain unused in the cell. Our expectation is that the assimilating protein will therefore be unlikely to retain this capability, as it evolves. Similarly, only a single extra link should be generated at the second-neighbor level, third-neighbor level, etc. Consistent with the evidence discussed above, the number of links created by assimilation is approximately independent of the total network size. Party hubs typically are centrally-located within modules, while date hubs often function to stitch together large-scale modules in the cell. It may be that duplication-only models are unrealistically fragmented (Table II) because their modules are not properly attached with date hubs; instead, the modules are disconnected components.

Sensitivity analysis

The DUNE model has four parameters. One parameter, the DU rate d , is estimated from empirical data. The other three are adjustable parameters: the NE rate μ , the divergence probability ϕ , and the assimilation probability a . To gain a better understanding of how these parameters affect the final network structure, starting with each organism's set of parameters, we systematically adjusted all 4 parameters. Results are shown in Figure 11.

As expected, the divergence parameter ϕ was positively correlated with both the modularity Q and the diameter. When $\phi \approx 0$, the network quickly reaches a fully-connected state, where all proteins are linked to all other proteins. Consequently, the network is not organized into modules, and all distances in the network are equal to 1. The 'thinning out' of duplicate links is therefore essential to generate non-trivial network features. The opposite occurs when the NE rate μ is too low: $f_1 \approx 0$, and the network evolves towards a completely disconnected state.

Why does gene duplication lead to modularity? Consider an initially uniform network. When a node is duplicated at random, this causes the original node, the copy, and their immediate neighbors to share more links internally than they do with the rest of the network. Subsequent duplications amplify this effect: if a node that has 10 links within a module but only 2 external links is duplicated, there are now 20 internal and 4 external links (prior to post-duplication divergence).

Interestingly, Q has a weak negative correlation with the assimilation parameter a . When a is large, the probability to link to distant neighbors of the target protein is relatively high, and mutated proteins have a non-

negligible chance to generate links to proteins outside of their target’s pathway. This causes modules to blur at the edges; their member proteins will share a higher number of links to other modules than for a low a network. Although the modularity is reduced for a high a network, it does not disappear entirely. Similarly, there is a sharp decrease in Q as the NE rate surpasses the DU rate, indicating the important role of the DU mechanism in modular organization.

Diameter is also negatively correlated with a . When a single NE event has a significant chance to generate links to the target protein’s neighbors, this tends to reduce the overall separation of proteins in the network.

DISCUSSION

The relevance of selection to PPI network evolution has been a topic of considerable debate [82], particularly in the context of higher-order network features, such as modularity. A number of authors have argued that specific selection programs are required to generate modular networks, such as oscillation between different evolutionary goals [83–88]. However, previous work has shown that gene duplication by itself, in the absence of both natural selection and neofunctionalization, can generate modular networks [89, 90]. Consistent with the findings of [89, 90], modularity in our model is primarily generated by gene duplications (Figure 11; see SI for sensitivity analysis). Unfortunately, duplication-only models err in their predictions of other network properties (Tables II and S2; Figure S7). A well-known problem with duplication models is that they generate excessively fragmented networks, with only about 20% of the proteins in the largest component. This is in sharp contrast to real PPI networks, which have 73% to 89% of their proteins in the largest component. Neofunctionalization-only models have most of their proteins in the largest component, but are significantly less modular than real networks. As shown in Table II, by modeling duplication and neofunctionalization simultaneously, the DUNE model generates networks which have the modularity found in duplication-only models, while retaining most proteins in the largest component. This lends support to the idea that gene duplication contributes to the modularity found in real biological networks, and that protein modules can arise under neutral evolution, without requiring complicated assumptions about selective pressures. This is consistent with recent experimental work characterizing a real-world fitness landscape, showing that it is primarily shaped by neutral evolution [91].

One model for the fate of duplicate genes is *subfunctionalization* (SF). In SF, the original and the duplicate genes are both free to lose their redundant functions, so they can evolve freely until they exactly reproduce the ancestral function [92]. The post-duplication divergence in our model is similar in spirit to SF, but it differs in two significant ways: (1) in our model, the link loss is

completely asymmetric, and (2) a fraction $(1 - \phi)$ of the redundant links are retained, so, unlike SF, not all of the redundancy is eliminated. For the first point, empirical evidence suggests that the divergence is asymmetric [31], although the assumption of *complete* asymmetry would likely need to be revisited to build a finer-grained model. Second, genetic regulatory networks have been shown to be robust to random link deletions, indicating that these networks retain some degree of redundancy [93–95]. *In silico* evidence suggests that a more accurate picture may be of a transient period of functional divergence, followed by prolonged neofunctionalization, resulting in only a partial loss of redundancy [96]. This is consistent with our model.

One example of a known biological mechanism which should lead to assimilation is domain shuffling, the copy-and-pasting of part of one protein into another [97, 98]. The neofunctionalization mechanism described here is quite general, and includes domain shuffling, among other methods of PPI creation. A PPI formed via domain shuffling will often be the result of a binding site duplication. Assuming the binding is due to simple surface similarity, the initial link will be to the protein which had its domain copied. The likelihood of binding to neighbors of the original protein should depend only on the probability that each interaction is due to surface similarity because the copied binding site will be identical to the original.

The role of domain shuffling in assimilation raises the question of whether domains should be modeled explicitly, rather than representing proteins as integral units. Previous work indicates that overall PPI network topology is robust to the details of domain shuffling [99]. Moreover, while proteins which have experienced domain shuffling have a higher average degree than other proteins, high- and low-degree proteins are equally likely to acquire new interactions this way [100]. Because the creation of new links by domain shuffling should be topologically very similar to the creation of new links by other neofunctionalization events, we believe our model is a reasonable implementation of this mechanism, as it applies to the evolution of network topology.

Previous estimates of NE rates in eukaryotes have varied widely, generally falling in the range of 100 to 1000 changes/genome/Myr [24, 46, 52], or on the order of 0.1 change/gene/Myr. However, more recent empirical work has identified several problems with the methods used to obtain these estimates, suggesting that *de novo* link creation is much less common than previously thought [48]. This is consistent with our model. The best-fit values of our NE rate μ are in the range of 10^{-5} to 10^{-4} /gene/Myr (Table I), which in all three organisms are considerably slower than the duplication rates d .

Biologically, many of the interactions created by our neofunctionalization mechanism are expected to initially be weak, non-functional interactions. The results of [101] suggest that strong functional interactions are correlated with hydrophobicity, which in turn is correlated with

promiscuity. We posit that initially weak, non-functional interactions are an essential feature of PPI evolution, as they provide the ‘raw material’ for the subsequent evolution of functional interactions. If this reasoning is correct, one consequence should be that hub proteins are, on average, more important to the cell than non-hub proteins. This has been found to be true: both degree [57] and betweenness centrality [68] have positive correlations with essentiality, indicating that hub proteins are often critical to the cell’s survival.

We have described here a model for how eukaryotic protein networks evolve. The model, called DUNE, implements two biological mechanisms: (1) gene duplications, leading to a superfluous copy of a protein that can change rapidly under new selective pressures, giving new relationships with other proteins and (2) a protein can undergo random mutations, leading to neofunctionalization, the *de novo* creation of new relationships with other proteins. Neofunctionalization can lead to assimilation, the formation of extra novel interactions with the other proteins in the target’s neighborhood. Biological evidence suggests that this type of mechanism exists. Our specific implementation is based on a simple geometric surface-compatibility argument for the observed transitivity in PPI networks. This is, of course, a heavily simplified model of PPI network evolution, and there are many biological factors which have not been included. However, our relatively simple model shows good agreement with 10 topological properties in yeast, fruit flies, and humans. One finding is that PPI networks can evolve modular structures, just from these random forces, in the absence of specific selection pressures. We also find that the most central proteins also tend to be the oldest. This suggests that looking at the structures of present-day protein networks can give insight into their evolutionary history.

Methods

Genome-wide PPI screens have a high level of noise [19], and specific interactions correlate poorly between data sets [54]. We found that several large-scale features differed substantially between types of high-throughput experiments (see SI). Due to concerns about the accuracy and precision of data obtained through high-throughput screens, we chose to work with ‘high-confidence’ data sets consisting only of pairwise interactions confirmed in small-scale experiments, which we downloaded from the public HitPredict database [102]. We found sufficient high-confidence data in yeast (*S. cerevisiae*), fruit flies (*D. melanogaster*), and humans (*H. sapiens*).

All simulations and network feature calculations were carried out in Matlab. Our scripts are freely available for download at <http://interacto.me>. We computed betweenness centralities, clustering coefficients, shortest paths, and component sizes using the MatlabBGL package. Modularity values were calculated with the algo-

rithm of [103]. All comparisons (except the degree distribution) are between the largest connected components of the simulated and experimental data.

Due to the human network’s somewhat larger size, most dynamical features were calculated once per 50 time steps for the human network, but were updated at every time step in the yeast and fly networks. For dynamical plots, the y coordinates of the trend line are medians-of-medians. The amount of time elapsed per time step (the x coordinate) varies between simulations. We binned the time coordinates to the nearest 10 million years for yeast and fly, and 25 million years for human. When multiple values from the same simulation fell within the same bin, we used the median value. We then calculated the median value between simulations. Scatter plot trend lines are calculated in a similar way. The trend line represents the median response variable (C , b , or ℓ) value over all nodes within a single simulation with degree k . The y coordinate of the trend line is therefore the median (across 50 simulations) of these median response variables. This median-of-medians includes all simulations that have nodes of a given degree.

Empirical data

We downloaded large-scale data sets from BioGRID [104], and used the Wilcoxon rank-sum test to compare aggregate statistical features across various experimental types in yeast (*S. cerevisiae*) and humans (*H. sapiens*) [78]. As expected, we found that data obtained by affinity capture was significantly different than pairwise experimental data (primarily yeast two-hybrid and *in vitro* complexation), as the affinity capture interactions represent entire complexes, which is somewhat different information than the pairwise interactions we are attempting to capture using our model. However, more surprisingly, the only feature to show significant agreement between pair-wise techniques was the eigenvalue distribution of the walk matrix ($P > 0.05$). Further sub-dividing the individual techniques into smaller data sets containing only results obtained in single experiments, we discovered that, again, only the spectra agreed between different screens.

Note that, due to the small size of the fly network, there may be too many missing links to obtain an accurate description the network’s large-scale topology. Although, by appropriate parameter tuning, our model is able to accurately reproduce the fly network, it is possible that different parameters will be required to match the fly network once it becomes more fully characterized experimentally. The data sets considered here do not include interactions which are enabled through post-translational modifications. Although these data sets are far from complete, and may be susceptible to false-positive detections, these appear to be the most accurate data available at the present time.

ACKNOWLEDGMENTS

We would like to thank Tom MacCarthy and Anna Panchenko for valuable discussions.

-
- [1] D. Hughes. Microbial genetics: Exploiting genomics, genetics and chemistry to combat antibiotic resistance. *Nature Reviews Genetics*, 4:432–441, 2003.
- [2] R.T. Cirz, J.K. Chin, D.R. Andes, V. de Crécy-Lagard, W.A. Craig, and F.E. Romesberg. Inhibition of mutation and combating the evolution of antibiotic resistance. *PLoS Biology*, 3:e176, 2005.
- [3] I.W. Taylor, R. Linding, D. Warde-Farley, Y. Liu, C. Pesquita, D. Faria, S. Bull, T. Pawson, Q. Morris, and J.L. Wrana. Dynamic modularity in protein interaction networks predicts breast cancer outcome. *Nature Biotechnology*, 27:199–204, 2009.
- [4] M. Lynch, M. O’Hely, B. Walsh, and A. Force. The probability of preservation of a newly arisen gene duplicate. *Genetics*, 159(4):1789–1804, 2001.
- [5] C.-T. Ting, S.-C. Tsaur, S. Sun, W.E. Browne, Y.-C. Chen, N.H. Patel, and C.-I. Wu. Gene duplication and speciation in *Drosophila*: evidence from the *Odysseus* locus. *Proc. Natl. Acad. Sci. USA*, 101(33):12232–12235, 2004.
- [6] J. Dutkowski and J. Tiuryn. Phylogeny-guided interaction mapping in seven eukaryotes. *BMC Bioinformatics*, 10:393, 2009.
- [7] T. Ito, K. Tashiro, S. Muta, R. Ozawa, T. Chiba, M. Nishizawa, K. Yamamoto, S. Kuhara, and Y. Sakaki. Toward a protein-protein interaction map of the budding yeast: A comprehensive system to examine two-hybrid interactions in all possible combinations between the yeast proteins. *Proc. Natl. Acad. Sci. USA*, 97:1143–1147, 2000.
- [8] P. Uetz, L. Giot, G. Cagney, T. Mansfield, R. Judson, J. Knight, D. Lockshon, V. Narayan, M. Srinivasan, P. Pochart, A. Qureshi-Emili, Y. Li, B. Godwin, D. Conover, T. Kalbfleisch, G. Vijayadamodar, M. Yang, M. Johnston, S. Fields, and J. Rothberg. A comprehensive analysis of protein-protein interactions in *Saccharomyces cerevisiae*. *Nature*, 403:623–627, 2000.
- [9] N.J. Krogan, G. Cagney, H. Yu, G. Zhong, X. Guo, A. Ignatchenko, J. Li, S. Pu, N. Datta, A.P. Tikuisis, T. Punna, J. Peregrin-Alvarez, M. Shales, X. Zhang, M. Davey, M.D. Robinson, A. Paccanaro, J.E. Bray, A. Sheung, B. Beattie, D.P. Richards, V. Canadien, A. Lalev, F. Mena, P. Wong, A. Starostine, M.M. Canete, J. Vlasblom, S. Wu, C. Orsi, S.R. Collins, S. Chandran, R. Haw, J.J. Rilstone, K. Gandhi, N.J. Thompson, G. Musso, P. St. Onge, S. Ghanny, M.H.Y. Lam, G. Butland, Atlaf Ul A.M., S. Kanaya, A. Shilatifard, E. O’Shea, J.S. Weissman, J.C. Ingles, T.R. Hughes, J. Parkinson, M. Gerstein, S.J. Wodak, A. Emili, and J.F. Greenblatt. Global landscape of protein complexes in the yeast *Saccharomyces cerevisiae*. *Nature*, 440:637–643, 2006.
- [10] H. Yu, P. Braun, M.A. Yildirim, I. Lemmens, K. Venkatesan, J. Sahalie, T. Hirozane-Kishikawa, F. Gebreab, N. Li, N. Simonis, T. Hao, J.F. Rual, A. Dricot, A. Vazquez, R.R. Murray, C. Simon, L. Tardivo, S. Tam, N. Svrikapa, C. Fan, A.S. de Smet, A. Motyl, M.E. Hudson, J. Park, X. Xin, M.E. Cusick, T. Moore, C. Boone, M. Snyder, F.P. Roth, A.L. Barabási, J. Tavernier, D.E. Hill, and M. Vidal. High-quality binary protein interaction map of the yeast interactome network. *Science*, 5898:104–110, 2008.
- [11] E. Marcotte, M. Pellegrini, H. Ng, D. Rice, T. Yeates, and D. Eisenberg. Detecting protein function and protein-protein interactions from genome sequences. *Nature*, 402:86–90, 1999.
- [12] M. Pellegrini, E. Marcotte, M. Thompson, D. Eisenberg, and T. Yeates. Assigning protein functions by comparative genome analysis: protein phylogenetic profiles. *Proc. Natl. Acad. Sci. USA*, 96:4285–4288, 1999.
- [13] A. Valencia and F. Pazos. Computational methods for the prediction of protein interactions. *Current Opinion in Structural Biology*, 12:368–373, 2002.
- [14] S. Gomez, W. Noble, and A. Rzhetsky. Learning to predict protein-protein interactions from protein sequences. *Bioinformatics*, 19:1875–1881, 2003.
- [15] R. Jothi, M. Kann, and T. Przytycka. Predicting protein-protein interaction by searching evolutionary tree automorphism space. *Bioinformatics*, 21:241–250, 2005.
- [16] Y. Liu, N. Liu, and H. Zhao. Inferring protein-protein interactions through high-throughput interaction data from diverse organisms. *Bioinformatics*, 21:3279–3285, 2005.
- [17] B. Shoemaker and A. Panchenko. Deciphering protein-protein interactions. part II. computational methods to predict protein and domain interaction partners. *PLoS Computational Biology*, 3:e43, 2007.
- [18] L. Burger and E. van Nimwegen. Accurate prediction of protein-protein interactions from sequence alignments using a Bayesian method. *Mol. Syst. Biol.*, 4:165, 2008.
- [19] C.M. Deane, Ł. Salwiński, I. Xenarios, and D. Eisenberg. Protein interactions: two methods for assessment of the reliability of high throughput observations. *Molecular & Cellular Proteomics*, 1(5):349–356, 2002.
- [20] T. Yamada and P. Bork. Evolution of biomolecular networks – lessons from metabolic and protein interactions. *Nature Reviews Molecular Cell Biology*, 10:791–803, 2009.
- [21] Susumu Ohno. *Evolution by Gene Duplication*. Springer-Verlag, New York, 1970.
- [22] J. Zhang. Evolution by gene duplication: an update. *Trends in Ecology & Evolution*, 18(6):292–298, 2003.
- [23] H. Xiao, N. Jiang, E. Schaffner, E.J. Stockinger, and E. van der Knaap. A retrotransposon-mediated gene duplication underlies morphological variation of tomato fruit. *Science*, 319:1527–1530, 2008.
- [24] A. Wagner. How the global structure of protein interaction networks evolves. *Proc. R. Soc. Lond. B*, 270:457–466, 2003.

- [25] A.L. Koch. Enzyme evolution. i. the importance of untranslatable intermediates. *Genetics*, 72(2):297–316, 1972.
- [26] J.S. Taylor and J. Raes. Duplication and divergence: the evolution of new genes and old ideas. *Annual Review of Genetics*, 38:615–643, 2004.
- [27] M. Lynch and J.S. Conery. The evolutionary fate and consequences of duplicate genes. *Science*, 290:1151–1155, 2000.
- [28] S. Maslov, K. Sneppen, and K.K. Eriksen, K.A. and Yan. Upstream plasticity and downstream robustness in evolution of molecular networks. *BMC Evolutionary Biology*, 4:9, 2004.
- [29] A. Vázquez, A. Flammini, A. Maritan, and A. Vespignani. Modeling of protein interaction networks. *ComplexUs*, 1:38–44, 2003.
- [30] M. Kellis, B.W. Birren, and E.S. Lander. Proof and evolutionary analysis of ancient genome duplication in the yeast *Saccharomyces cerevisiae*. *Nature*, 428:617–624, 2004.
- [31] X. Gu, Z. Zhang, and W. Huang. Rapid evolution of expression and regulatory divergences after yeast gene duplication. *Proc. Natl. Acad. Sci. USA*, 102:707–712, 2005.
- [32] M. Lynch and B. Walsh. *Genetics and Analysis of Quantitative Traits*. Sinauer, Sunderland, MA, 1998.
- [33] I. Ispolatov, A. Yuryev, I. Mazo, and S. Maslov. Binding properties and evolution of homodimers in protein-protein interaction networks. *Nucleic Acids Res.*, 33(11):3629–3635, 2005.
- [34] S. Jones and J.M. Thornton. Principles of protein-protein interactions. *Proc. Natl. Acad. Sci. USA*, 93:13–20, 1996.
- [35] A. Tovchigrechko and I.A. Vakser. How common is the funnel-like energy landscape in protein-protein interactions? *Protein Science*, 10(8):1572–1583, 2001.
- [36] F. Wu, F. Towfic, D. Dobbs, and V. Honavar. Analysis of protein-protein dimeric interfaces. In *IEEE International Conference on Bioinformatics and Biomedicine*. Fremont, CA, 2007.
- [37] A.E. Vinogradov and O.V. Anatskaya. Loss of protein interactions and regulatory divergence in yeast whole-genome duplicates. *Genomics*, 93(6):534–542, 2009.
- [38] H. Escriva, S. Bertrand, P. Germain, M. Robinson-Rechavi, M. Umbhauer, J. Cartry, M. Duffraisse, L. Holland, H. Gronemeyer, and V. Laudet. Neofunctionalization in vertebrates: The example of retinoic acid receptors. *PLoS Genetics*, 2(7):e102, 2006.
- [39] D.T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, 81:2340–2361, 1977.
- [40] Z. Gu, A. Cavalcanti, F.C. Chen, P. Bouman, and W.H. Li. Extent of gene duplication in the genomes of *Drosophila*, nematode, and yeast. *Molecular Biology and Evolution*, 19:256–262, 2002.
- [41] L. Gao and H. Innan. Very low gene duplication rate in the yeast genome. *Science*, 306(5700):1367–1370, 2004.
- [42] M. Lynch and J.S. Conery. The evolutionary demography of duplicate genes. *Journal of Structural and Functional Genomics*, 3:35–44, 2003.
- [43] N. Osada and H. Innan. Duplication and gene conversion in the *Drosophila melanogaster* genome. *PLoS Genetics*, 4(12):e1000305, 2008.
- [44] J.A. Cotton and R.D.M. Page. Rates and patterns of gene duplication and loss in the human genome. *Proceedings of the Royal Society B*, 272, 2005.
- [45] D. Pan and L. Zhang. Quantifying the major mechanisms of recent gene duplications in the human and mouse genomes: a novel strategy to estimate gene duplication rates. *Genome Biol.*, 8(8):R158, 2007.
- [46] P. Beltrao and L. Serrano. Specificity and evolvability in eukaryotic protein interaction networks. *PLoS Computational Biology*, 3, 2007.
- [47] M. Lynch, W. Sung, K. Morris, N. Coffey, C.R. Landry, E.B. Dopman, W.J. Dickinson, K. Okamoto, S. Kulkaarni, D.L. Hartl, and W.K. Thomas. A genome-wide view of the spectrum of spontaneous mutations in yeast. *Proc. Natl. Acad. Sci. USA*, 105(27):9272–9277, 2008.
- [48] T.A. Gibson and D.S. Goldberg. Questioning the ubiquity of neofunctionalization. *PLoS Computational Biology*, 5:e1000252, 2009.
- [49] M. Kimura. Some problems of stochastic processes in genetics. *Ann. Math. Stat.*, 28:882–901, 1957.
- [50] J.B. Walsh. How often do duplicated genes evolve new functions? *Genetics*, 139:421–428, 1995.
- [51] X. He and J. Zhang. Rapid subfunctionalization accompanied by prolonged and substantial neofunctionalization in duplicate gene evolution. *Genetics*, 169:1157–1164, 2005.
- [52] J. Berg, M. Lassig, and A. Wagner. Structure and evolution of protein interaction networks: a statistical model for link dynamics and gene duplications. *BMC Evolutionary Biology*, 4:51–63, 2004.
- [53] N. Pržulj, D.G. Corneil, and I. Jurisica. Modeling interactome: scale-free or geometric? *Bioinformatics*, 20:3508–3515, 2004.
- [54] E.J. Deeds, O. Ashenberg, and E.I. Shakhnovich. A simple physical model for scaling in protein-protein interaction networks. *Proc. Natl. Acad. Sci. USA*, 103:311–316, 2006.
- [55] P. Erdős and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.*, 5:17–61, 1960.
- [56] A. Patil, K. Kinoshita, and H. Nakamura. Domain distribution and intrinsic disorder in hubs in the human protein-protein interaction network. *Protein Science*, 19:1461–1468, 2010.
- [57] H. Jeong, S.P. Mason, A.L. Barabási, and Z.N. Oltvai. Lethality and centrality in protein networks. *Nature*, 411:41–42, 2001.
- [58] H.M. Ku, T. Vision, J. Liu, and S.D. Tanksley. Comparing sequenced segments of the tomato and *Arabidopsis* genomes: large-scale duplication followed by selective gene loss creates a network of synteny. *Proc. Natl. Acad. Sci. USA*, 97(16):9121–9126, 2000.
- [59] J. Travers and S. Milgram. An experimental study of the small world problem. *Sociometry*, 32(4):425–443, 1969.
- [60] J. Leskovec and E. Horvitz. In *Proceedings of the 17th international conference on World Wide Web*. ACM, New York, 2008.
- [61] S.H. Yook, Z.N. Oltvai, and A.L. Barabási. Functional and topological characterization of protein interaction networks. *Proteomics*, 4:928–942, 2004.
- [62] M.E.J. Newman and M. Girvan. Finding and evaluating community structure in networks. *Phys. Rev. E*, 69:026113, 2004.
- [63] M.E.J. Newman. Modularity and community structure

- in networks. *Proc. Natl. Acad. Sci. USA*, 103:8577–8582, 2006.
- [64] Sergei Maslov and Kim Sneppen. Specificity and Stability in Topology of Protein Networks. *Science*, 296(5569):910–913, 2002.
- [65] C. von Mering, R. Krause, B. Snel, M. Cornell, S.G. Oliver, S. Fields, and P. Bork. Comparative assessment of large-scale data sets of protein-protein interactions. *Nature*, 31:399–403, 2002.
- [66] E. Ravasz, A.L. Somera, D.A. Mongru, Z.N. Oltvai, and A.L. Barabási. Hierarchical organization of modularity in metabolic networks. *Science*, 297:1551–1555, 2002.
- [67] A.L. Barabási and Z.N. Oltvai. Network biology: understanding the cell’s functional organization. *Nature Reviews Genetics*, 5:101–113, 2004.
- [68] M.P. Joy, A. Brock, D.E. Ingber, and S. Huang. High-betweenness proteins in the yeast protein interaction network. *Journal of Biomedicine and Biotechnology*, 2:96–103, 2005.
- [69] D. Zhao, Z. Liu, and J. Wang. Duplication: a mechanism producing disassortative mixing networks in biology. *Chin. Phys. Lett.*, 24(10):2766, 2007.
- [70] X. Wan, S. Cai, J. Zhou, and Z. Liu. Emergence of modularity and disassortativity in protein-protein interaction networks. *Chaos*, 20:045113, 2010.
- [71] P. Aloy and R.B. Russell. Potential artefacts in protein-interaction networks. *FEBS Lett.*, 530:253–254, 2002.
- [72] L. Hakes, J.W. Pinney, D.L. Robertson, and S.C. Lovell. Protein-protein interaction networks and biology - what’s the connection? *Nature Biotechnol.*, 26:69–72, 2008.
- [73] C.R. Woese. Bacterial evolution. *Microbiol. Rev.*, 51(2):221–271, 1987.
- [74] P.L. Krapivsky and S. Redner. Organization of growing random networks. *Physical Review E*, 63:066123, 2001.
- [75] H. Qin, H.H.S. Lu, W.B. Wu, and W.H. Li. Evolution of the yeast protein interaction network. *Proc. Natl. Acad. Sci. USA*, 100(22):12820–12824, 2003.
- [76] G. Zhu, H. Yang, R. Yang, J. Ren, B. Li, and Y.C. Lai. Uncovering evolutionary ages of nodes in complex networks. *Eur. Phys. Jour. B.*, 85(3):106, 2012.
- [77] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graphs over time: densification laws, shrinking diameters and possible explanations. *Knowledge Discovery in Databases: PKDD 2005*, 2005.
- [78] J.D. Han, N. Bertin, T. Hao, D.S. Goldberg, G.F. Berriz, L.V. Zhang, D. Dupuy, A.J. Walhout, M.E. Cusick, F.P. Roth, and M. Vidal. Evidence for dynamically organized modularity in the yeast protein-protein interaction network. *Nature*, 430(6995):88–93, 2004.
- [79] P.M. Kim, L.J. Lu, Y. Xia, and M.B. Gerstein. Relating three-dimensional structures to protein networks provides evolutionary insights. *Science*, 314(5807):1938–1941, 2006.
- [80] D. Ekman, S. Light, A.K. Bjorklund, and A. Elofsson. What properties characterize the hub proteins of the protein-protein interaction network of *Saccharomyces cerevisiae*? *Genome Biol.*, 7(6):R45, 2006.
- [81] G.P. Singh, M. Ganapathi, and D. Dash. Role of intrinsic disorder in transient interactions of hub proteins. *Proteins*, 66(4):761–765, 2007.
- [82] M. Lynch. The evolution of genetic networks by non-adaptive processes. *Nat. Rev. Genet.*, 8(10):803–813, 2007.
- [83] H. Lipson, J.B. Pollack, and N.P. Suh. On the origin of modular variation. *Evolution*, 56(8):1549–1556, 2002.
- [84] N. Kashtan and U. Alon. Spontaneous evolution of modularity and network motifs. *Proc. Natl. Acad. Sci. USA*, 102(9):13773–13778, 2005.
- [85] K. Komurov and M. White. Revealing static and dynamic modular architecture of the eukaryotic protein interaction network. *Molecular Systems Biology*, 3:110, 2007.
- [86] G.P. Wagner, M. Pavlicev, and J.M. Cheverud. The road to modularity. *Nature Reviews Genetics*, 8:921–931, 2007.
- [87] C. Espinosa-Soto and A. Wagner. Specialization can drive the evolution of modularity. *PLoS Computational Biology*, 6(3):e1000719, 2010.
- [88] O.S. Soyer. Fate of a duplicate in a network context. In K. Dittmar and D. Liberles, editors, *Evolution After Gene Duplication*, pages 215–228. Wiley-Blackwell, 2010.
- [89] J. Hallinan. Gene duplication and hierarchical modularity in intracellular interaction networks. *BioSystems*, 74:51–62, 2004.
- [90] R.V. Solé and S. Valverde. Spontaneous emergence of modularity in cellular networks. *J. R. Soc. Interface*, 5(18):129–133, 2008.
- [91] R.T. Hietpas, J.D. Jensen, and D.N.A. Bolon. Experimental evolution of a fitness landscape. *Proc. Natl. Acad. Sci. USA*, 108(19):7896–7901, 2011.
- [92] M. Lynch and A. Force. The probability of duplicate gene preservation by subfunctionalization. *Genetics*, 154:459–473, 2000.
- [93] M.L. Siegal and A. Bergman. Waddington’s canalization revisited: developmental stability and evolution. *Proc. Natl. Acad. Sci. USA*, 99(16):10528–10532, 2002.
- [94] A. Bergman and M.L. Siegal. Evolutionary capacitance as a general feature of complex gene networks. *Nature*, 424(6948):549–552, 2003.
- [95] A. Wagner and J. Wright. Alternative routes and mutational robustness in complex regulatory networks. *BioSystems*, 88:163–172, 2007.
- [96] T. MacCarthy and A. Bergman. The limits of subfunctionalization. *BMC Evolutionary Biology*, 7:213, 2007.
- [97] J.V. Moran, R.J. DeBardinis, and H.H. Kazazian. Exon shuffling by L1 retrotransposition. *Science*, 283:1530–1534, 1999.
- [98] L. Patthy. Genome evolution and the evolution of exon-shuffling – a review. *Gene*, 238:103–114, 1999.
- [99] K. Evlampiev and H. Isambert. Modeling protein network evolution under genome duplication and domain shuffling. *BMC Syst. Biol.*, 1:49, 2007.
- [100] D.B. Cancherini, G.S. Franca, and S.J. de Souza. The role of exon shuffling in shaping protein-protein interaction networks. *BMC Genomics*, 11(Suppl 5):S11, 2010.
- [101] M. Heo, S. Maslov, and E. Shakhnovich. Topology of protein interaction network shapes protein abundances and strengths of their functional and nonspecific interactions. *Proc. Natl. Acad. Sci. USA*, 108(10):4258–4263, 2011.
- [102] A. Patil, K. Nakai, and H. Nakamura. Hitpredict: a database of quality assessed protein-protein interactions in nine species. *Nucleic Acids Research*, 39 (suppl 1):D744–D749, 2011.
- [103] V.D. Blondel, J.-L. Guillaume, R. Lambiotte, and

- E. Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008:P10008+, 2008.
- [104] C. Stark, B.J. Breitkreutz, T. Reguly, L. Boucher, A. Breitkreutz, and M. Tyers. BioGRID: a general repository for interaction datasets. *Nucleic Acids Res.*, 34:D535–9, 2006.
- [105] S. Scherer. *A Short Guide to the Human Genome*. Cold Spring Harbor, NY, 2008.
- [106] K.P. Byrne and K.H. Wolfe. The yeast gene order browser: combining curated homology and syntenic context reveals gene fate in polyploid species. *Genome Res.*, 15:1456–1461, 2005.
- [107] J.W. Drake, B. Charlesworth, D. Charlesworth, and J.F. Crow. Rates of spontaneous mutation. *Genetics*, 148:1667–1686, 1998.
- [108] Neil A. Campbell, Jane B. Reece, and Lawrence G. Mitchell. *Biology*. Benjamin/Cummings, Menlo Park, CA, 5th edition, 1999.
- [109] A. Clauset, C.R. Shalizi, and M.E.J. Newman. Power-law distributions in empirical data. *SIAM Review*, 51:661–703, 2009.
- [110] J. Shlens. A tutorial on principal component analysis. 2009.
- [111] R. Pastor-Satorras, E. Smith, and R.V. Solé. Evolving protein interaction networks through gene duplication. *Journal of Theoretical Biology*, 222:199–210, 2003.
- [112] M. Middendorf, E. Ziv, and C.H. Wiggins. Inferring network mechanisms: the *Drosophila melanogaster* protein interaction network. *Proc. Natl. Acad. Sci. USA*, 102(9):3192–3197, 2005.
- [113] M. O’Connor and M. Lawrence. Judgmental forecasting and the use of available information. In G. Wright and P. Goodwin, editors, *Forecasting with Judgment*. Wiley, Chichester, 1998.
- [114] M. Hibon and S. Makridakis. The M-3 competition: results, conclusions, and implications. *International Journal of Forecasting*, 16:461–476, 2000.

SUPPORTING INFORMATION

Randomness conjecture

We describe here a test of our randomness conjecture, $q = 1/N$. The implication is that the average number of NE links per protein should be independent of N . Another possibility is that q is constant, implying that the number of NE links per protein is proportional to N . To test this, we compared NE rates in the human and yeast PPI networks. The total number of proteins in humans is estimated to be 22740 [105] and in yeast, 5616 [106]. The number of mutations to coding DNA is approximately 0.004/genome/replication in humans and 0.0027/genome/replication in yeast [107]. If the number of new links is proportional to N , then, based on the number of proteins and the mutation rate, there should be roughly 600% more links created by NE in humans than in yeast. However, by counting the number of nonredundant interactions in duplicate gene pairs, it has been shown empirically that the average number of links cre-

ated by NE per protein is only about 8% higher in humans than in yeast [51]. These results support the conjecture that the probability for a protein to receive a new link via point mutation is approximately independent of N , as previously noted in [46]. Due to the finite copy number of proteins, as well as the compartmentalization of eukaryotic cells, we regard it as unlikely that proteins will simultaneously acquire multiple links to targets in different locations in the cell, or which are involved in divergent biological processes.

Large-scale duplications

Our model does not explicitly consider simultaneous duplication of multiple genes (chromosomal duplications, whole genome duplications, etc.). However, as shown in Table I, duplication rates in our model are considerably higher than neofunctionalization rates, so that, on average, there are multiple duplications per neofunctionalization event. Sequential duplications of this type may be thought of as an (imperfect) representation of multi-gene duplications. The advantage of this approach is that we do not require separate rates for each duplication scale (one could imagine an extremely detailed model which included separate rates for gene duplication, gene-pair duplication, gene-triplet duplication, etc.). The downside is that our implementation of larger-scale duplications will generally include some genes which have been duplicated multiple times, and others which have not been duplicated at all. A potential mitigating factor is that the rate of gene loss (and evolution in general) following genome duplication is very high [30, 58], so even a completely faithful large-scale duplication would likely be altered within short order.




FIG. S1. **Walk matrix eigenvalues.** Shown are eigenvalue (λ) distributions in human (green), yeast (blue), and fly (red). Heavy lines are the median values from 50 simulations, and light lines are results of individual simulations.

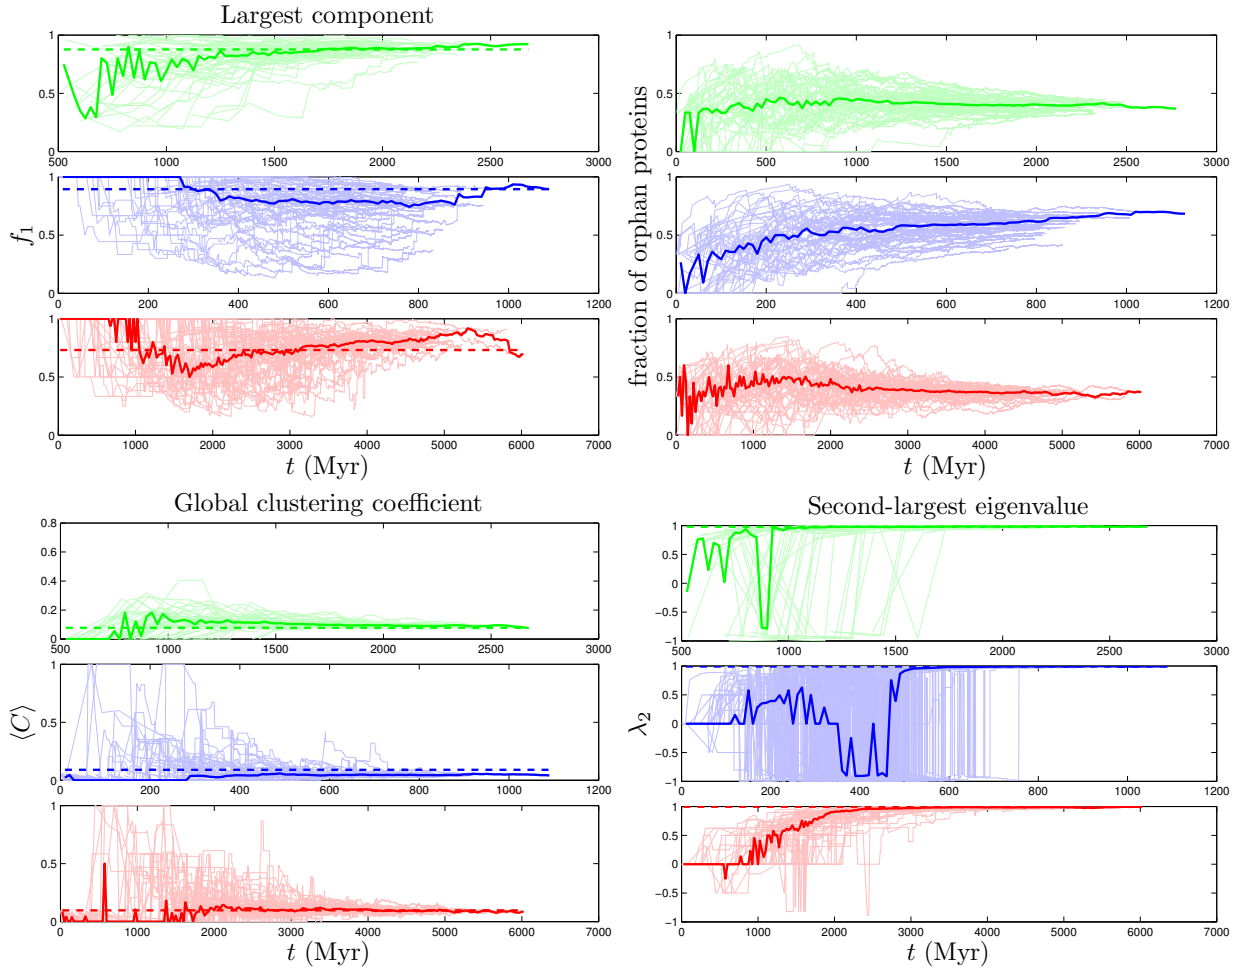


FIG. S2. **Dynamical features.** Shown are the evolution of (A) the largest component size, (B) the fraction of orphan proteins, (C) the global clustering coefficient, and (D) the second-largest eigenvalue of the walk matrix, in human (green), yeast (blue), and fly (red). Light lines indicate the evolutionary trajectories of 50 individual simulations, and the heavy line is the median value. Empirical data values are shown as a dashed line, where available.

Eigenvalues

The connectivity of a network can be expressed by its *adjacency matrix*, an $N \times N$ matrix \mathbf{A} , in which the entries A_{ij} equal 1 if a link exists between proteins i and j , and 0 otherwise. If \mathbf{A} is normalized by column, then the entries describe the rates of a transition from i to j in one time step. The distribution of eigenvalues $p(\lambda)$ is called the network's *spectrum* (Figure S1). This matrix and its eigenvalues can be interpreted in terms of a process in which a random walker starts on one node i , and, over a series of time steps, reaches another node j . Intuitively, this can be thought of as a signal propagation rate: if one protein is affected by an external signal, how long does it take that signal to diffuse through the network? The eigenvalues λ of this 'walk matrix' describe the rate at which a random walk on the network reaches steady-state. The second-largest eigenvalue (λ_2) determines the rate of convergence of the random walk (Figure S2). A

larger value of λ_2 indicates a slower signal propagation rate.

Error tolerance

We measured the 'error tolerance' as described in [29]: we examined the decrease of f_1 when nodes (and their accompanying edges) were deleted from the network either (1) at random, or (2) according to their degree, starting with the most well-connected node. Results for the simulated and experimental networks were very similar (Figure S4).

Simulation length: early versus late evolution

The total time elapsed during our simulations varies considerably, with yeast and human simulations running

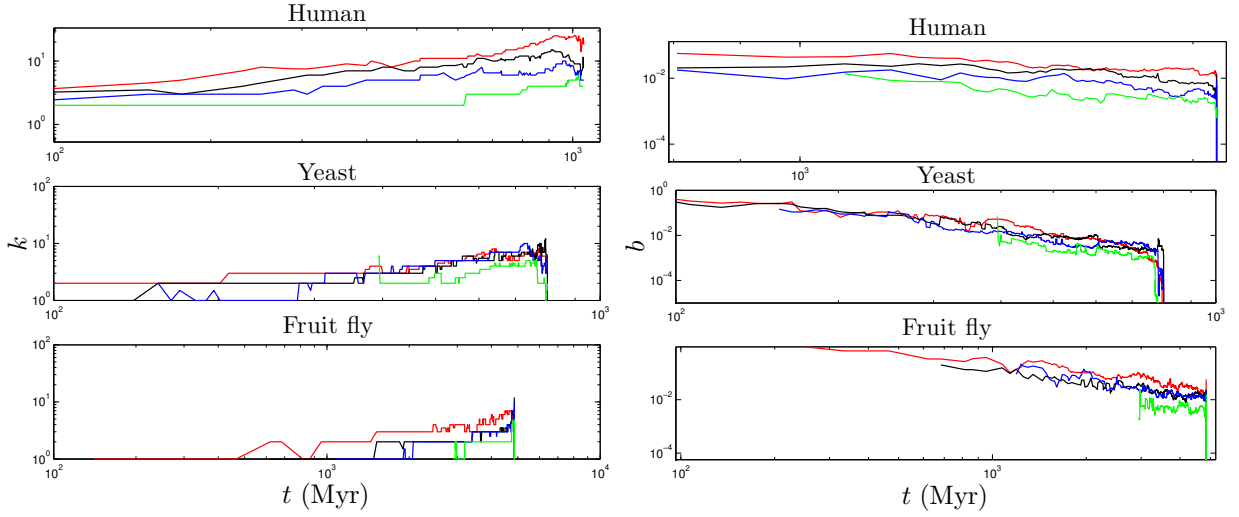


FIG. S3. **Individual protein centrality scores.** Evolution of degree (A) and betweenness (B) for proteins introduced to the network at different times in humans (top), yeast (middle), and flies (bottom). The 1st protein (one of the two initial proteins) is shown in red, the 6th protein in black, the 11th protein in blue, and the 101st protein in green. Curves are median values from 50 simulations.

about 1 to 2 billion years, and the fly simulations about 5 billion years. This is compared to the rough estimate of 3.5 billion years since the origin of life on Earth [108]. The exceptionally long duration of the fly simulations are due to the very low gene duplication rate ($d = 0.001/\text{gene}/\text{Myr}$). The aim of our model is to describe the evolution of PPI networks with all their present-day machinery. Gene duplication, in the form in which it exists today, certainly would not have existed at the origin of life! The initial state in our model con-

sists of two interacting proteins. Biologically, these are two polypeptides (or, more likely, RNA molecules) in a pre-biotic soup, that happen to interact in a way that is mutually beneficial. Each of these molecules has the ability to replicate. This autonomous replication of individual proteins corresponds to ‘gene duplication’ in the very early stages of evolution. However, this is a very different conceptual underpinning for the duplication mechanism, and it seems unlikely to share the present-day values of the duplication rate. Because, in the early stages of evolution, each time step represents a very long duration in real time, it is likely that this accounts for the discrepancy in total time elapsed.




FIG. S4. **Error tolerance.** Shown are error tolerance curves in human (green), yeast (blue), and fly (red). Circles indicate proteins deleted randomly, and squares indicate proteins deleted starting with the most well-connected protein and removing proteins in descending order.

Fitting functions

The degree distribution obeys a power law in its tail, $p(k) \sim k^{-\gamma}$ [109], with $\gamma \approx 3$ (Table S1), implying that hub proteins are more common than would be expected for a randomly connected network, which would have an exponentially decaying $p(k)$. The closeness distribution

	γ	β	ξ	α	δ
Yeast	2.8(2)	2.4(1)	1.8(2)	1.2(2)	0.32(6)
Fly	3.1(2)	2.2(4)	0.8(5)	0.8(7)	0.0(3)
Human	2.8(1)	2.3(1)	1.5(3)	1.3(2)	0.0(1)

TABLE S1. **Scaling exponents.** Distributional exponents ($p(k) \sim k^{-\gamma}$, $p(b) \sim b^{-\beta}$) were estimated using the maximum likelihood method of [109]. Other exponents ($C \sim k^{-\xi}$, $b \sim k^{\alpha}$, $n \sim k^{-\delta}$) were estimated using nonlinear regression. Due to the relatively small sizes of the data sets, there is considerable uncertainty in these estimates.

$p(\ell)$ is approximately Gaussian, with mean 0.17 and standard deviation 0.03 in humans, mean 0.19 and standard deviation 0.03 in yeast, and mean 0.13 and standard deviation 0.03 in flies. Closeness is a measure of distance, indicating that the distances within the network are essentially a random walk in ‘node space’. The betweenness distribution also follows a power law in its tail. This is an indication of modular structure, due to the overrepresentation of ‘bridge’ proteins, relative to a randomly connected network.

All species examined show a power law decay in clustering coefficient as a function of degree, $C \sim k^{-\xi}$. Poorly-connected proteins therefore tend to have *higher* clustering coefficients, meaning that a greater fraction of their neighbors are mutually connected.

Disassortative mixing was quantified for the yeast PPI network in [64] as a power law *decrease* in median neighbor degree, $n \sim k^{-\delta}$. This is consistent with our data, although the very small estimated value of $\delta = 0.32$ indicates only a slight negative relation (Table S1). Interestingly, $\delta = 0$ in both human and fly networks, indicating that disassortativity may be a trait unique to the yeast network.

Principal component analysis

We examined six features which calculate a value for each node in the network: degree centrality, clustering coefficients, closeness centrality, eigenvalue spectrum, betweenness centrality, and mean nearest-neighbor degree. To quantify the independence of these features, we used principal component analysis (PCA) [110]. Each feature assigns a value to each node in the network, giving a $6 \times N$ data matrix, where each row represents a feature (signal), and each column is a node (sample). We subtract the mean and divide by the standard deviation of each row. This results in a standardized data matrix, denoted by \mathbf{Y} . The 6×6 correlation matrix for each species is defined as $\mathbf{C} \equiv \frac{1}{N-1} \mathbf{Y} \mathbf{Y}^T$:

$$\mathbf{C}_h = \begin{bmatrix} 1 & 0.10 & 0.87 & 0.56 & 0.02 & -0.12 \\ 0.10 & 1 & -0.04 & 0.09 & 0.01 & 0.03 \\ 0.87 & -0.04 & 1 & 0.44 & -0.01 & -0.08 \\ 0.56 & 0.09 & 0.44 & 1 & 0.00 & 0.34 \\ 0.02 & 0.01 & -0.01 & 0.00 & 1 & -0.02 \\ -0.12 & 0.03 & -0.08 & 0.34 & -0.02 & 1 \end{bmatrix}, \quad (4)$$

$$\mathbf{C}_y = \begin{bmatrix} 1 & 0.03 & 0.91 & 0.43 & -0.02 & -0.21 \\ 0.03 & 1 & -0.06 & 0.04 & 0.01 & 0.03 \\ 0.91 & -0.06 & 1 & 0.39 & -0.01 & -0.14 \\ 0.43 & 0.04 & 0.39 & 1 & -0.03 & 0.34 \\ -0.02 & 0.01 & -0.01 & -0.03 & 1 & -0.04 \\ -0.21 & 0.03 & -0.14 & 0.34 & -0.04 & 1 \end{bmatrix}, \quad (5)$$




FIG. S5. **Laplacian eigenvector participation.** Elements of the eigenvector of the Laplacian matrix (defined as $\mathbf{K} - \mathbf{A}$, where \mathbf{K} is a diagonal matrix with the degree of node i as element K_{ii}) associated with the largest eigenvalue vs. protein age index (time of introduction) in the yeast simulation. Details of this method are discussed in [76]. Heavy lines are the median values from 50 simulations, and light lines are results of individual simulations. The inset plot shows the trend line with a rescaled y -axis.

$$\mathbf{C}_f = \begin{bmatrix} 1 & 0.15 & 0.62 & 0.36 & -0.11 & -0.15 \\ 0.15 & 1 & -0.08 & 0.06 & -0.10 & -0.02 \\ 0.62 & -0.08 & 1 & 0.40 & 0.02 & -0.16 \\ 0.36 & 0.06 & 0.40 & 1 & 0.00 & 0.30 \\ -0.11 & -0.10 & 0.02 & 0.00 & 1 & -0.05 \\ -0.15 & -0.02 & -0.16 & 0.30 & -0.05 & 1 \end{bmatrix}. \quad (6)$$

The entries of each \mathbf{C} are (from left-to-right, and top-to-bottom): degree centrality, clustering coefficients, betweenness centrality, closeness centrality, eigenvalue spectrum, and mean nearest-neighbor degree. Many of the off-diagonal elements of the \mathbf{C} matrices are close to zero, suggesting that the features are to a large extent independent of one another.

To perform PCA, we diagonalized each correlation matrix,

$$\mathbf{C} = \mathbf{S} \mathbf{\Lambda} \mathbf{S}^T, \quad (7)$$

where $\mathbf{\Lambda}$ is a diagonal matrix of eigenvalues and \mathbf{S} has the eigenvectors of \mathbf{C} as its columns. As shown in Figure S6, the degree and betweenness show similar loadings on the first two principal components, reflecting the nearly linear relation between these centrality scores (Figure 6).




FIG. S6. **Principal component analysis.** Shown are the factor loadings and scores on the first two principal components. Data scores are shown in red, and blue lines represent feature loadings.

The eigenvalue matrices Λ are given by:

$$\Lambda_h = \begin{bmatrix} 2.27 & & & & & \\ & 1.23 & & & & \\ & & 1.02 & & & \\ & & & 0.98 & & \\ & & & & 0.40 & \\ & & & & & 0.11 \end{bmatrix}, \quad (8)$$

$$\Lambda_y = \begin{bmatrix} 2.20 & & & & & \\ & 1.30 & & & & \\ & & 1.01 & & & \\ & & & 0.98 & & \\ & & & & 0.43 & \\ & & & & & 0.08 \end{bmatrix}, \quad (9)$$

$$\Lambda_f = \begin{bmatrix} 1.95 & & & & & \\ & 1.24 & & & & \\ & & 1.13 & & & \\ & & & 0.91 & & \\ & & & & 0.44 & \\ & & & & & 0.32 \end{bmatrix}. \quad (10)$$

(Zeros have been suppressed for clarity.) The fraction of variance explained by the i th principal component is given by $\Lambda_{ii} / \sum_j \Lambda_{jj}$. The closer the number of components required to explain most of the variance is to the total number of input signals, the more independent the signals are. In yeast and humans, 4 components are required to explain 90% of the variance; in fruit flies, it requires 5 components. Linear transformations are able to only modestly reduce the dimensionality of the problem, suggesting that each feature contributes unique information about the network's structure. This does not, of course, rule out the possibility of the existence of other independent, informative features, a far more complicated question which is outside the scope of this current work.

Comparison to other models

Our model is rooted in previous modeling efforts. The basic framework for our model combines the gene duplication mechanism described in [29] with a link creation mechanism inspired by [52]. The principal difference between our model and previous models is that our model considers duplication and mutation simultaneously. The previous models we examined attempted to construct the PPI network from a single mechanism. Another significant difference is our assimilation mechanism. To the best of our knowledge, previous work has not explicitly modeled proteins integrating into biological pathways.

We compare the DUNE model to four models previously proposed for PPI networks. Two were evolutionary models: (1) the Vázquez model of DU followed by rapid loss-of-function mutations [29] and (2) the Berg 'link dynamics' model of point mutations coupled with a PA-like 'rich-get-richer' rule for assigning new interactions [52]. (A slightly different DU model is presented by Pastor-Satorras [111]. However, because the Vázquez model has been shown to be a better fit to experimental data [112], we have limited our DU-only comparison here to the Vázquez model.) Two others were static models (models of present-day networks that do not simulate the network's evolutionary path) that consider the primary organizing principle to be nonspecific interactions between proteins: (1) random geometric (RG), a mathematical model where proteins are randomly scattered in a 2 to 4 dimensional box, and any proteins close enough to one another form an interaction [53] and (2) the 'MpK' desolvation model, which assigns interactions based on proteins' exposed hydrophobic surface areas [54]. For reference, we also calculated results for an Erdős-Rényi (ER) random graph with N and $\langle k \rangle$ set by the data [55].

These models were originally validated against different features of the empirical network, making it difficult to directly compare them. To characterize these models in greater detail, we coded each of these models, and ran 50 simulations of each model with identical parameters and starting conditions. Using Matlab, we coded the

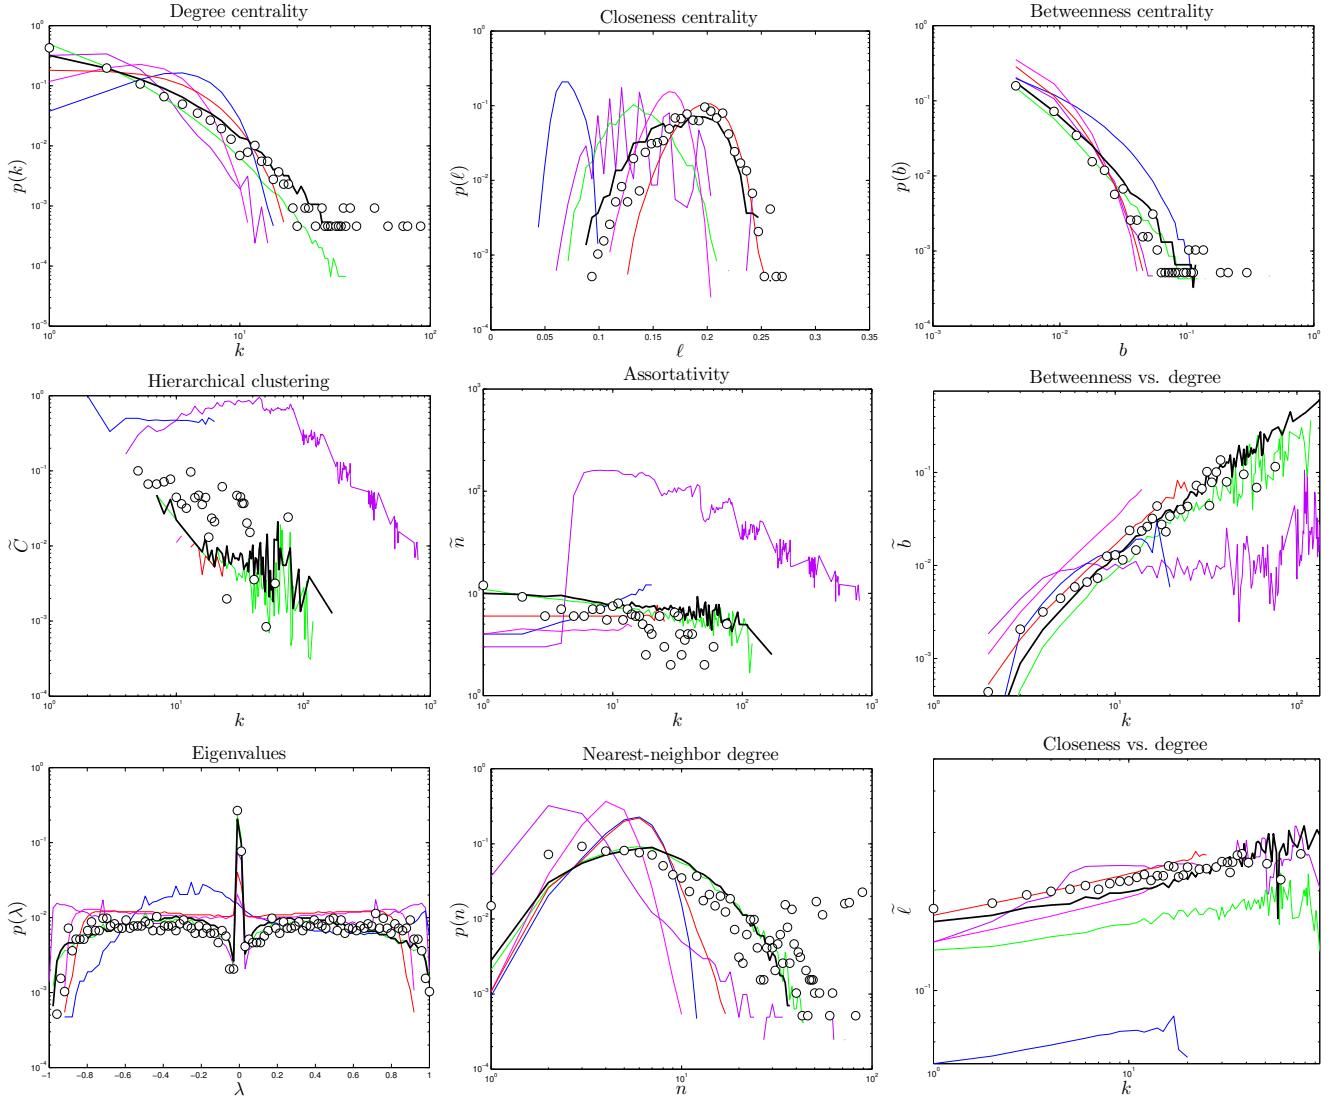


FIG. S7. **Model comparison.** Comparison of five other models to the yeast PPI network: Vázquez [29] (green), Berg [52] (red), random geometric [53] (dark blue), MpK desolvation [54] (purple), and ER random graph [55] (brown). For reference, DUNE model results are shown as a black line. Dots represent high-confidence experimental yeast data, and solid lines are median values over 50 simulations.

Vázquez [29], Berg [52], RG [53], and MpK [54] models as described in the original papers. Since each model was originally parametrized for older yeast PPI data sets, we re-optimized the parameters for our yeast data as follows. We used a Monte Carlo simulation to adjust each model's parameters to minimize the total symmetric mean absolute percentage error values (SMAPE; see below) for the yeast HitPredict data set.

For the Vázquez model, we used a value of 0.582 for the post-duplication divergence probability, and a value of 0.083 for the dimerization probability. As noted by previous authors, duplication-only simulations produce networks which are extremely fragmented [89]. We observed that the Vazquez simulations typically had around 20% of their nodes in the largest connected component (Ta-

ble II). Since most of the network features we examined are limited to the largest component, in order to make a reasonable comparison of the Vazquez simulation results to the data, we allowed the simulated network to grow until its number of links met or exceeded 5 times the number of links in the data, $K \geq 5K_{\text{data}}$. Since the largest component is not always exactly 20% of the total nodes, this stopping condition is somewhat arbitrary; however, results for this model seem robust to small changes in the stopping condition. For the Berg model, we used the empirically estimated duplication rate of 0.01/gene/Myr, and found best-fit values of 24.5/gene/Myr for the mutation rate, and $N_{\text{data}} - 98$ proteins for the initial network size. For the RG model, we used a $45.5 \times 45.5 \times 45.5$ 'box' with a maximum interaction radius of 3.92. For

	$p(k)$	$p(\ell)$	$p(b)$	$C(k)$	$n(k)$	$b(k)$	$p(\lambda)$	E.T.	E.T. (k)	$p(n)$	$\ell(k)$	Total
DUNE	0.47	0.37	0.36	0.60	0.29	0.24	0.14	0.10	0.20	0.56	0.14	3.48
Vázquez	0.52	0.81	0.39	0.62	0.27	0.34	0.12	0.10	0.15	0.50	0.26	4.11
Berg	0.70	0.45	0.72	0.67	0.15	0.19	0.29	0.12	0.30	0.86	0.06	4.51
RG	0.81	0.98	0.61	0.83	0.22	0.23	0.33	0.31	0.45	0.89	0.49	6.15
MpK	0.82	0.79	0.68	0.89	0.88	0.58	0.18	0.18	0.18	0.82	0.16	6.18
ER	0.83	0.70	0.78	0.53	0.22	0.41	0.28	0.07	0.32	0.93	0.07	5.14

TABLE S2. Symmetric mean absolute percentage error (SMAPE) of simulation versus experiment in yeast (Eq. 11). ‘E.T.’ is the error tolerance curve with random protein removal, and ‘E.T. (k)’ is the error tolerance curve with highest-degree proteins removed first. ‘DUNE’ is the model described here, ‘Vázquez’ is the DU-only model of [29], ‘Berg’ is the link dynamics model [52], ‘RG’ is random geometric [53], ‘MpK’ is the physical desolvation model presented in [54], and ‘ER’ is an Erdős-Rényi random graph [55]. For each comparison, the lowest value is shown in bold.

the MpK model, the number of exposed surface residues was 19, the fraction of exposed hydrophobic residues was $M = 0.230 \pm 0.110$ (mean \pm standard deviation), and the best-fit linear equation relating M to the binding threshold was $1.09M + 1.04$.

The Vázquez simulations were initialized with 2 connected nodes, and the simulation was allowed to run until $K \geq 5K_{\text{data}}$. The Berg simulations were initialized with $N_{\text{data}} - 98$ randomly connected nodes, then run until $N = N_{\text{data}}$. The RG and MpK models (which are not evolutionary models and therefore create the network all at once) were set up as described in the original papers.

To characterize the networks, we computed several network properties:

Single-value:: modularity Q , diameter D , fraction of nodes in the largest component f_1 , global clustering coefficient $\langle C \rangle$, and average protein degree in the largest component $\langle k \rangle$ (Table II)

Distributional:: degree $p(k)$, betweenness $p(b)$, closeness $p(\ell)$, eigenvalue $p(\lambda)$, and nearest-neighbor degree $p(n)$ distributions

Scatter plot:: closeness vs. degree $\ell(k)$, clustering coefficient vs. degree $C(k)$, betweenness vs. degree $b(k)$, median nearest-neighbor degree vs. degree $n(k)$, and error tolerance curves

and compared these features to those of empirical data from yeast. As shown in Figure S7, we found that none

of the previous models capture the full set of network properties.

To quantify agreement with the data for non-single-value features, we calculated the symmetric mean absolute percentage error (SMAPE) between simulation and experiment [113, 114]:

$$\text{SMAPE} = \frac{1}{Y} \sum_i^Y \frac{|y_i - y_i^{\text{data}}|}{y_i + y_i^{\text{data}}}, \quad (11)$$

where Y is the number of data points, and y_i and y_i^{data} denote the i th point of the response variable (Table S2) in the simulated and experimental data, respectively. For the distributional features, Y is the number of bins (arbitrarily chosen to be 100) minus the number of bins in which $y_i + y_i^{\text{data}} = 0$. For non-distributional (scatter plot) features, Y is the number of k values with values for both simulation and experiment. There are many possible measures of accuracy (such as the widely-used root mean squared error); we used SMAPE for two reasons. First, because it relies on absolute value, SMAPE does not over-emphasize the impact of outliers. Second, dividing by $y_i + y_i^{\text{data}}$ ensures that the magnitude of the response variable does not overwhelm the sum. This is significant for the non-distributional features. For example, in a plot of betweenness vs. degree (Figure 6), we are just as interested in the overlap of the low-betweenness, low-degree region of the curve as we are with the high-betweenness, high-degree region. SMAPE values are collected in Table S2. As shown in Tables II and S2, while previous models accurately reproduce certain features of the PPI network, only the DUNE model provides a reasonable across-the-board fit.